arXiv:1611.09123v1 [cond-mat.soft] 28 Nov 2016

Many-Body Coarse-Grained Interactions using Gaussian Approximation Potentials

S.T. John

Peterhouse, University of Cambridge

September 2016

This dissertation is submitted for the degree of Doctor of Philosophy.

Summary

Coarse-graining (CG) is an approach to molecular simulations in which molecules are described at a lower resolution than individual atoms, thereby reducing the number of interaction sites and speeding up the simulation. Considering the formal connection between all-atom and CG models given by statistical mechanics, a CG model is deemed to be *consistent* with the all-atom model when the CG interactions are described by the many-body potential of mean force (PMF).

In practice, all current CG methods only consider a limited set of interaction terms between a very small number of CG sites; non-bonded interactions are typically described by pairwise radial interactions between sites. Effective interactions can be obtained by matching certain properties of the all-atom model, such as pairwise distance distributions or forces. However, this results in a lack of *representability*: CG models that only include these limited site-based interactions cannot represent both structural and thermodynamic properties of the underlying all-atom system.

In the field of atomistic simulations, many-body interactions have been successfully described using the Gaussian Approximation Potential (GAP), a machine-learning approach based on Gaussian process regression for obtaining flexible, high-dimensional interaction potentials. The present work transfers GAP to coarse-graining (GAP-CG), in which the CG interactions are derived from observations of forces in the all-atom model. This includes the development of solutions to specific challenges such as inhomogeneous length scales and insufficient sampling. GAP-CG allows us to approximate the PMF using *molecular* terms that correspond to pairs and triplets of entire molecules.

In applying GAP-CG to methanol and benzene clusters and bulk systems, this thesis demonstrates that this approach can recover the PMF, allowing a representation of both structural (including orientational) distributions *and* forces. Furthermore, this thesis studies how the accuracy of force predictions depends on the hyperparameters.

Many-body interactions are computationally expensive. At the example of water-solvated benzene, this thesis shows that a solvent-free CG model using GAP interactions can be faster than the corresponding all-atom model while being significantly more accurate than current approaches to coarse-graining.

Dedicated to my father
who would have liked to read this

Contents

1. Introduction **5**

1.1. Molecular Simulations 5
1.2. Coarse-Graining 8
1.3. Many-Body Potentials 11
1.4. Outline 13

2. Classical Molecular Dynamics **15**

2.1. Equations of Motion 16
2.2. Force Fields 17
2.3. Constant-Temperature Simulations and Thermostats 19
2.4. Free Energy 20
2.4.1. Collective Variables and the Potential of Mean Force 21
2.4.2. Mean Forces 23
2.5. Summary 24

3. Coarse-Graining Approaches **25**

3.1. Top-down *vs.* Bottom-up Models 26
3.2. Mapping from All-Atom to Coarse-Grained Model 29
3.3. Consistency and the Many-Body Potential of Mean Force 31
3.3.1. Momentum Space 32
3.3.2. Configuration Space 33
3.4. Structure-based Approaches 34
3.4.1. Direct Boltzmann Inversion 35
3.4.2. Iterative Boltzmann Inversion 36
3.4.3. Inverse Monte Carlo 38
3.4.4. Relative Entropy 39
3.5. Force-based Approaches 40
3.5.1. Instantaneous Collective Forces 41
3.5.2. Practical Approach: MSCG/FM 42
3.5.3. Connection to Relative Entropy 43
3.6. Summary 44

4. Gaussian Approximation Potential **45**
4.1. Gaussian Process Regression . . . 47
4.1.1. Bayesian Modelling . . . 47
4.1.2. Distributions over Functions . . . 48
4.1.3. Covariance Function and Hyperparameters . . . 51
4.1.4. Predictions . . . 53
4.1.5. Sums and Derivatives . . . 58
4.1.6. Sparsification . . . 60
4.2. Interaction Potentials . . . 61
4.2.1. Locality Approximation . . . 62
4.2.2. Descriptors . . . 65
4.2.3. Core Potential . . . 69
4.2.4. Multiple Descriptors . . . 70
4.2.5. Selection of Sparse Points . . . 71
4.3. Application to Coarse-Graining: GAP-CG . . . 74
4.3.1. Noise . . . 76
4.3.2. Inhomogeneous Length Scales . . . 77
4.3.3. Sampling Issues . . . 79
4.4. Summary . . . 81

5. GAP-CG: Recovering the Potential of Mean Force **83**
5.1. Model Systems: All-Atom Description and CG Mapping . . . 85
5.1.1. Methanol . . . 85
5.1.2. Benzene . . . 86
5.1.3. Simulation Protocol . . . 87
5.2. Methanol Clusters . . . 88
5.2.1. Monomer . . . 88
5.2.2. Dimer . . . 93
5.2.3. Trimer . . . 106
5.2.4. Tetramer . . . 117
5.3. Bulk Simulations . . . 125
5.3.1. Simulation Protocol . . . 126
5.3.2. Locality Test . . . 127
5.3.3. Methanol . . . 130
5.3.4. Benzene . . . 139
5.4. Influence of Hyperparameters, Sparsification, and Training Data . . . 145
5.4.1. Hyperparameters . . . 145
5.4.2. Sparsification . . . 147
5.4.3. Noisy Training Data . . . 148
5.5. Summary . . . 151

6. Practical Application: Solvent-Free Coarse-Graining **153**
6.1. Water-Solvated Benzene . . . 154
6.1.1. All-Atom Description . . . 154
6.1.2. Benzene Dimer in Different Environments . . . 156
6.1.3. CG Potentials . . . 160
6.1.4. MD Results . . . 163
6.2. Recovering a Pair PMF from the CG Potential . . . 171
6.3. Simulation Speed . . . 172
6.4. Summary . . . 174

7. Conclusions **175**
7.1. Scaling . . . 177
7.2. Limitations of GAP-CG . . . 178
7.3. Outlook and Future Work . . . 179

A. Appendix **183**
A.1. GAP hyperparameters . . . 183
A.2. Pair potentials . . . 185
A.3. Length scale hyperparameter and transfer function . . . 189
A.4. Mean forces for benzene dimer in different environments . . . 190

Bibliography **193**

Abbreviations

QM	Quantum Mechanics
MD	Molecular Dynamics
MM	Molecular Mechanics
AA	All-Atom
CG	Coarse-Graining
GAP	Gaussian Approximation Potential
CV	Collective Variable
PMF	Potential of Mean Force
MF	Mean Force
ICF	Instantaneous Collective Force
DBI	Direct Boltzmann Inversion
IBI	Iterative Boltzmann Inversion
IMC	Inverse Monte Carlo
MSCG/FM	Multiscale Coarse-Graining/Force Matching
COM	Centre Of Mass
RDF	Radial Distribution Function
ADF	Angular Distribution Function
RMSE	Root Mean Squared Error

1. Introduction

1.1. Molecular Simulations

Understanding processes on an atomic or molecular scale is important in a wide range of fields from materials science to molecular biology. This includes diverse topics, from studying the formation and propagation of cracks or elucidating the properties of novel "nano" materials to investigating the behaviour of proteins, lipids, and nucleic acids or studying the interactions between small molecules and binding sites to aid in the development of new drugs. Ultimately, only experiments can give definite, authoritative answers to questions about these processes.

However, there are many questions that cannot directly be answered by experiments alone. This can be for a number of reasons. We do not always have the tools available to access all the properties in which we are interested. Experiments might only provide limited resolution without being able to access microscopic details of structure or dynamics. The experimental probe might require a modification of, or otherwise interfere with, the system we want to study, possibly changing its behaviour. Moreover, experiments tend to be very expensive and require a large amount of direct involvement and human time. Consequently, the number of experiments performed is often small, resulting in comparatively large errors and uncertainty, especially in biological studies.

Computer simulations of molecular systems can help us fill these gaps. They allow us to access *all* properties of the system within the simulation, and let us see the effect of manipulating any part of the system. This extends to manipulations that would be difficult or impossible in the real world (for example, we can simply replace one type of molecule with another). In approaches that simulate the dynamics of a system, we can also study its *dynamic* behaviour and its response to virtual probes (such as stretching a molecule to determine its force–extension curve).

It is important to note that simulations do not make experiments redundant. A simulation is just a model, and its accuracy and its predictions, ultimately, need to be validated by experiments [1]. In fact, computer simulations can often be complementary to experiments, providing a bridge between real world and theory. Simulations can

assist us in devising and refining experiments, and can help us understand experimental observations [2].

To be useful, a model should accurately represent those features of the real system in which we are interested. Moreover, it should make quantitative predictions that can be compared, where possible, with experimental results. Only then is a simulation a valuable tool.

All-atom simulations

The most accurate theory for describing processes on an atomistic scale is quantum mechanics (QM). By solving the time-dependent Schrödinger equation for a molecular system, we can obtain the time evolution of the probability distributions for the positions of electrons and nuclei. However, this is rarely feasible except for the smallest systems. Instead, we usually consider approximations to the full quantum-mechanical description of a system. The nuclei are much heavier than the electrons; thus, in many cases we can consider their motion separately (Born–Oppenheimer approximation [3]) and often the motion of the nuclei is described well by classical mechanics (Newton's laws of motion). This is the basis of Molecular Dynamics (MD), one of the main approaches to molecular simulations, which calculates the dynamic trajectories of discrete particles based on the forces acting on them [4]. We will explain this approach in more detail in Chapter 2.

For an MD simulation of atoms as classical particles at the positions of the nuclei, we require the forces as a function of their configuration. Using the Hellman–Feynman theorem [5], we can calculate the forces from a quantum-mechanical description of the electrons. Common computational approaches for this are Quantum Monte Carlo [6] or Density Functional Theory (DFT) [7]. Quantum Monte Carlo provides high accuracy, but also has a high computational cost that limits it to comparatively small systems containing on the order of a thousand electrons. DFT considers only the electron *density*, and can be considered the state-of-the-art approach for routine QM calculations. While DFT codes can now compute forces in systems containing tens of thousands of atoms [8], quantum-mechanical calculations are still too expensive for simulations of the dynamics of systems on biologically relevant scales.

In practice, molecular simulations are usually based on empirical force fields that approximate the true quantum-mechanical forces by a sum of independent interaction terms. Each term only depends on the positions of a few atoms within a cutoff and generally has a fixed functional form (see Section 2.2). This approach is referred to as

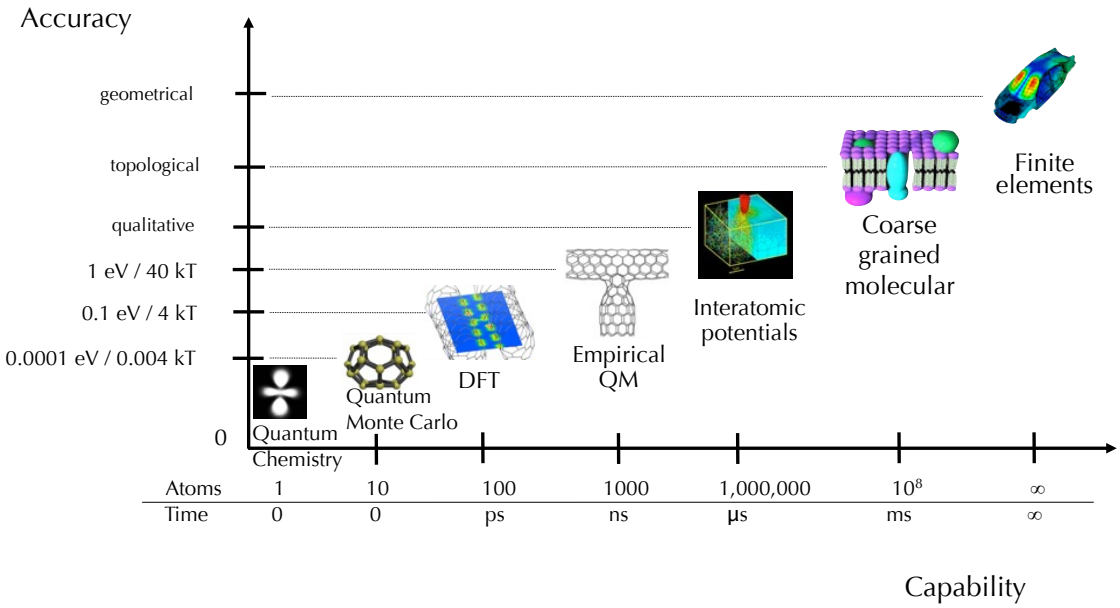

Figure 1.1: Scales of atomic and molecular simulations. *Figure courtesy of Gábor Csányi.*

"molecular mechanics" (MM). (We note that it is also possible to perform combined QM/MM simulations that treat certain parts of a model quantum-mechanically and use MM approximations for the remainder of the system.) The scales and accuracies of the different approaches are illustrated in Fig. 1.1.

Molecular mechanics is routinely used to simulate molecular systems. On commodity machines with typically 16 cores, it is possible to simulate systems containing several thousand atoms, reaching ratios of model time to computing time that are on the order of 10 ns/day.

Using highly specialised computers [9], simulations based on molecular mechanics have been extended up to the millisecond time scale for small systems such as ubiquitin with approximately ten thousand atoms in the solvated model [10]; on 512 nodes this achieved a performance of 15 µs/day. On general-purpose high-performance computing clusters, one of the largest all-atom MD simulations to date was of the HIV-1 capsid [11], with 64 million atoms in the solvated system. Obtaining dynamic trajectories of 0.1 µs required two weeks of computing time on 128 000 cores, corresponding to an average of 7 ns/day. A recent large-scale simulation of the influenza A virion with 5 million particles achieved a 5 µs trajectory [12] (no performance details published). These systems span length scales of up to a few hundred nanometres. We note that

such simulations are very expensive and cannot be repeated routinely.

The human drive to study larger and larger systems far outpaces the development of high-performance computing. Systems of interest in soft matter typically span micrometres, contain hundreds of millions of atoms and show behaviour on much longer time scales than is accessible in routine simulations [13]. Even with modern supercomputers, we are far away from being able to simulate systems such as an entire living cell.

Limitations of all-atom simulations

Accessing large systems remains a computational challenge. The number of particles in the simulation grows with the volume and hence cubically with the length scale. Moreover, with an increase in the length scale the *time* scale of the relevant motions also increases: On larger scales, it takes longer for the effect of a perturbation to ripple across the system. Furthermore, entropic effects become more important; these are generally softer and hence result in slower motions (consider, for example, the large-scale fluctuations of a lipid membrane *vs.* the small-scale fluctuations of individual lipid molecules). This means that not only do we have to simulate more particles, but also we have to carry out the simulations for longer.

The reach of all-atom simulations is limited by the available computational resources. The differential equations governing the motions of the atoms must be discretised, and the integration time step needs to be small enough to represent the fastest motion in the system (usually vibrations of bonds involving hydrogen atoms, if present in the system, due to their small mass; this is generally the case in biomolecular systems). For each time step, we need to calculate the forces on all the atoms based on their current positions. This is the bottleneck in molecular simulations.

1.2. Coarse-Graining

Accessing larger scales

In many cases we do not care about the fluctuations on the scale of individual atoms, but are interested in the behaviour of a system on a larger scale. In these cases we can extend the reach of our simulations to larger systems and longer time scales by employing lower-resolution models. This can be achieved by grouping together several atoms and representing them as single particles. This process is referred to as "coarse-graining" (CG), and the particles are usually referred to as "CG sites" [14].

We note that speeding up simulations is not the only application of coarse-graining. Another motivation is being able to only keep those degrees of freedom which (we think) are relevant for certain processes or behaviours, and to tweak the interactions at that scale, thereby allowing us to study the effects of eliminating degrees of freedom or changing the interactions at different scales.

Coarse-graining can take place on a wide range of scales. On the smallest scale, there are "united-atom" models in which heavy atoms are grouped with their bonded hydrogen atoms and described by new effective atom types, while all other atoms keep their identity. On the largest scale, a single CG site can represent several residues or molecules at once [15, 16].

Why coarse-graining speeds up simulations

Coarse-graining of a system leads to faster simulations for several reasons [17]:

First, coarse-graining reduces the number of particles that have to be simulated. Hence, the number of interactions that have to be computed can be much smaller, decreasing the computational effort required for each time step. This is an especially prominent advantage if we are only interested in the behaviour of a solute, such as a protein solvated by water. For this we can consider a "solvent-free" CG model in which we do not include the solvent molecules: their influence is only captured implicitly through the effective coarse-grained interactions between solutes, vastly reducing the number of interactions that need to be calculated in each time step. Moreover, longer-range interactions such as electrostatics are usually screened by the solvent, and hence the effective interactions between CG sites can have a much shorter range, further reducing the number of interactions that need to be calculated.

Second, the small-scale fluctuations on the atomic scale are no longer represented in the coarse-grained model, and high-frequency motions, such as covalent bond vibrations, are filtered out. Hence, CG simulations can employ a larger time step.

Third, fluctuations on the atomistic scale give rise to friction; for example, the atoms of two molecules may "interlock" and be unable to slide past each other immediately. Hence the CG model, without these fluctuations, evolves faster in time than the all-atom model. However, this actually changes and distorts the dynamic behaviour of the system. In general, one would like to retain this friction and its effect on the dynamics. (We note that in some cases, such as searching for stable states, the distortion of the dynamics may be irrelevant.) How to obtain accurate dynamics in a CG model is an area of ongoing research; this is the challenge of "real dynamics".

Coarse-graining is especially useful if models at different scales can be integrated [18]. An example of this would be a simulation of the active binding site of a protein with atomistic resolution, while the rest of the system is modelled at a much coarser level.

Like any simulation, coarse-graining has to be *accurate* to be useful. The CG simulation needs to reproduce the relevant properties of the system. This is a relative measure, and depends on the properties in which one is interested.

Limitations of coarse-graining

Coarse-graining faces three major challenges:

Transferability. Coarse-graining can be considered an average over fluctuations on the atomic scale, arriving at an effective CG interaction (this will be discussed in more detail in Chapter 3). This makes the CG interactions dependent on the state point (temperature, density, pressure). CG interactions parametrised for one state point cannot be assumed to reproduce the properties of the system at another, different state point. Hence, a given CG model is specific to the state point for which its interactions have been parametrised, and it cannot be assumed to be applicable in other contexts. The transferability of a CG model across state points needs to be confirmed empirically [19] or modelled explicitly in the interactions between CG sites [20–22].

Real dynamics. As noted previously, fluctuations on the atomic scale create friction between atoms and molecules. By coarse-graining, we average over these fluctuations. This means that the CG particles experience less friction, and motions in the CG model occur in less simulation time than in the all-atom model (corresponding to a scaling of the time axis). This could be seen as an advantage, as it provides even larger speedups. However, the faster time evolution is an artefact of artificial dynamics. Even if the CG model accurately reproduces structural distributions, the relation between "CG time" and real time is generally unknown, and cannot be assumed to be linear. To get an accurate coarse-grained simulation of the *dynamics* of a system, we have to account for the missing friction and for the impact of small-scale fluctuations by explicitly including drag and stochastic forces in the CG model [23].

Representability. To date, coarse-grained models generally assume that the CG interactions can be separated into a sum of simple terms, similar to those of the molecular mechanics force fields used in all-atom simulations. Specifically, interactions between

non-bonded CG sites are commonly assumed to be pairwise additive and radial (isotropic) "pair potentials". Based on this assumption, there are many different approaches that aim to derive effective pair potentials for a CG model, optimising for different properties of the system, such as structures, forces, or thermodynamics.

However, the interactions in the CG model are *effective* interactions resulting from an average over many atom–atom interactions. This leads to correlations between CG sites that, in general, cannot be decomposed into pairwise interactions.* In Chapter 3 we will discuss the connection between all-atom and coarse-grained models within a statistical mechanics framework. This leads to a definition of "consistency" for CG models based on matching probability distributions of the atomistic ensemble. A consistent CG model will reproduce all ensemble averages of an all-atom model that can be described in terms of the CG coordinates. To obtain a consistent CG model for N CG sites we need, in principle, an N-body potential to describe the CG interactions. It has been shown that it is impossible to reproduce both structural and thermodynamic properties of a system using only pair potentials [25].

Hence, CG models that rely on pair potentials necessarily trade off structural against thermodynamic accuracy. This is known as the representability problem: pair potentials cannot represent the "true" CG interactions. This issue will be investigated in this thesis, with the aim of overcoming this challenge.

1.3. Many-Body Potentials

A full N-body potential is generally infeasible because its computational cost scales exponentially with the number of sites. To obtain interactions that can be used in practice (and to realise CG simulations that are faster than the all-atom model), we need to approximate the N-body potential by separating it into a sum of simpler interaction terms. This leads to a trade-off between simplicity of the terms and accuracy of the resulting overall potential.

Whereas a coarse-grained model based on pair potentials suffers from the representability problem and is limited in the accuracy it can provide, this could be alleviated by considering more complex interactions that go beyond pair potentials. Both all-atom and CG models typically include angle and dihedral terms, which are three- and four-body interactions, respectively. However, these terms are usually limited to bonded short-range interactions with simple functional forms. They do not constitute a general

*In fact, even interactions between atoms can never be exactly represented by pair potentials [24].

many-body potential. In this thesis we focus on the non-bonded interactions that, so far, are predominantly described by pair potentials alone.

In the context of coarse-graining, many-body terms for non-bonded interactions have generally been neglected. They are often considered to be too computationally expensive and, as the primary aim of coarse-graining is a speedup of the simulation, only the simplest and cheapest interactions tend to be considered. Thus there has been only limited previous research into non-bonded many-body interactions in coarse-graining. There have been a few studies using three-body terms [26, 27]; however, these studies only considered one-site models of homogeneous liquids. Moreover, interaction terms are generally assumed to be a function of independent sites: when a molecule is described by multiple sites, this neglects the correlation between the sites describing a single molecule.

In this thesis we consider coarse-grained models for molecules that retain some internal structure and are described by more than one CG site, and we base the coarse-grained interactions on *molecular* terms that take into account all CG sites belonging to a molecule. This leads to more complex interactions that cannot be described by simple functional forms. To capture the interactions, we transfer the Gaussian Approximation Potential (GAP) approach to coarse-graining. Originally developed for atomistic modelling, GAP [28–30] is a very general approach that is based on Gaussian process regression [31, 32] to fit a flexible interaction potential to observations of energies and forces, without imposing a specific functional form on the potential.

We will show that we can successfully employ GAP to generate CG interactions based on forces sampled from the atomistic model, whilst still reproducing its full structure (as measured by a variety of structural distributions and correlations). This demonstrates that we can indeed create *consistent* CG models. Although GAP models are computationally expensive, for solvent-free CG models, which are relevant in many biomolecular simulations, these models can still provide a significant speedup compared to all-atom simulations, while being much more accurate than current approaches could hope to be.

1.4. Outline

In the first part of this thesis, we give an overview of classical molecular dynamics with a focus on all-atom models (Chapter 2) and review current approaches to coarse-graining (Chapter 3). In Chapter 4 we introduce Gaussian process regression (Section 4.1) and discuss the locality approximation and the descriptors that take into account our knowledge about the properties of physical systems and enable the Gaussian Approximation Potential (GAP) to provide effective and flexible many-body interactions (Section 4.2). In Section 4.3 we show how GAP can be used as a new, many-body approach to coarse-graining (GAP-CG). Inferring interactions only from force observations is a novel application of GAP, and we develop solutions to the challenges this brings up. In Chapter 5 we present the application of GAP-CG to methanol and benzene clusters and bulk systems. We show that the flexible molecular terms that are available in the GAP approach can capture atomistically consistent CG interactions, allowing a representation of both complex structural distributions *and* forces, and hence thermodynamic properties. Furthermore, we study how the accuracy of force predictions depends on the hyperparameters. In Chapter 6 we discuss solvent-free CG models and demonstrate for the example of water-solvated benzene that GAP-CG can have clear computational benefits while being significantly more accurate than current approaches to coarse-graining. In Chapter 7 we discuss our results, put them into the context of performance and scaling of CG models, and give an outlook to future work.

2. Classical Molecular Dynamics

The atomistic and coarse-grained simulations we consider in this thesis are based on Molecular Dynamics (MD). In MD, we model our system in terms of discrete (typically structureless, point-like*) particles.

Each particle, indexed by the subscript i, has a mass† m_i and is described by its position $\boldsymbol{r}_i$ and velocity $\boldsymbol{v}_i = \mathrm{d}\boldsymbol{r}_i/\mathrm{d}t$. The time evolution of positions and velocities is given by the equations of motion, which are integrated numerically; this will be discussed in Section 2.1. We thereby obtain a dynamic trajectory of the system, allowing calculations of both structural and dynamic properties.

The behaviour of the system depends on the interactions between the particles, defined in terms of the forces acting on them. The collection of the interaction terms is called the "force field" and will be described in Section 2.2.

In practice, many systems of interest are not isolated, but are in thermal equilibrium with the environment. In statistical mechanics, this is described by the canonical ensemble. To sample the canonical ensemble, the simulation needs to include a thermostat; this will be discussed in Section 2.3.

In thermal equilibrium, the system exchanges energy with is environment, and hence the internal energy (kinetic + potential energy) is no longer a constant of motion. Instead, the *free* energy becomes important to characterise the state of the system. Generally, we are interested in the (conditional) free energy as a function of one or more "collective variables" that characterise the configuration of the system on a larger scale; this is also referred to as the Potential of Mean Force (PMF). The negative gradients of the conditional free energy with respect to collective variables are *mean* forces. Free energy/PMF, collective variables and mean forces will be discussed in Section 2.4.

*In principle, simulations can also be carried out with anisotropic particles or rigid bodies, but this leads to more complicated equations, and point-like particles are most commonly used.

†In atomistic simulations, this is usually the average mass across the isotopes of the atom.

2.1. Equations of Motion

In atomistic simulations, each atom is modelled by a point-like particle at the position of the nucleus. The electrons are only modelled implicitly through the interactions they mediate between nuclei. Due to the large masses of the nuclei, quantum effects can generally be neglected, and their motion can be described by classical mechanics. This also holds for coarse-grained simulations, where the particle masses are even larger. The dynamics are then governed by Newton's laws of motion. The acceleration $\boldsymbol{a}_i$ of particle i is given by

$$m_i \boldsymbol{a}_i = m_i \frac{\mathrm{d}\boldsymbol{v}_i}{\mathrm{d}t} = m_i \frac{\mathrm{d}^2 \boldsymbol{r}_i}{\mathrm{d}t^2} = \boldsymbol{f}_i\big(\{\boldsymbol{r}_j\}\big), \tag{2.1.1}$$

where $\boldsymbol{f}_i$ is the force acting on the particle. In general, the force can be a function of the positions of all particles. When the forces are known, this differential equation can be integrated to obtain the positions of all the atoms as a function of time.

In general, the integration needs to be done numerically, by discretising Eq. (2.1.1). The most commonly used symplectic (phase space-conserving) integrator is the Velocity Verlet algorithm [33]. The new position and velocity after a time step Δt are calculated as

$$\begin{aligned} \boldsymbol{r}(t + \Delta t) &= \boldsymbol{r}(t) + \boldsymbol{v}(t)\,\Delta t + \frac{\boldsymbol{a}(t)}{2}\,(\Delta t)^2, \\ \boldsymbol{v}(t + \Delta t) &= \boldsymbol{v}(t) + \frac{\boldsymbol{a}(t) + \boldsymbol{a}(t + \Delta t)}{2}\,\Delta t, \end{aligned} \tag{2.1.2}$$

where for clarity we suppressed the index i.

The time step needs to be small enough to be able to describe the fastest motions in the system. Otherwise, discretisation errors would cause instabilities. In atomistic models, the fastest motion is generally the vibration of bonds with hydrogen atoms, if present in the system, due to their small mass. This requires a time step $\Delta t \lessapprox 1\,\mathrm{fs}$ [34].

Equation (2.1.2) describes Hamiltonian dynamics. When the forces are conservative, this corresponds to a constant-energy simulation of the microcanonical ensemble (also called the NVE ensemble, as the number of particles N, the volume V, and the energy E are fixed). To simulate systems at a constant temperature T, a thermostat algorithm (see Section 2.3) can be used to amend these equations to sample the canonical (or NVT) ensemble, though their general structure remains the same.

2.2. Force Fields

The equations of motion (2.1.2) depend on the forces acting on all the particles. These are determined by the force field, which gives the forces as a function of the configuration of the system. We generally assume conservative forces, which can be written as the negative gradient of the potential energy U,

$$\boldsymbol{f}_i = -\nabla_i U(\{\boldsymbol{r}_j\}),$$

where ∇_i is the gradient operator with respect to the position $\boldsymbol{r}_i$ of particle i. This applies to both all-atom and coarse-grained models. We now focus on all-atom simulations; the interactions in coarse-grained models will be discussed in Chapter 3.

In principle, the potential U can be a general function of the coordinates of all the atoms in the system. For example, this could be the quantum-mechanical (QM) total energy, obtained from (approximately) solving the QM Hamiltonian (for instance, using Density Functional Theory, Quantum Monte Carlo, or quantum chemistry methods such as Coupled Cluster Theory).

For simulations of larger systems such as biomolecules or complex materials, this would not be computationally feasible. Hence, it is often assumed that the potential energy of the system can be separated into a sum of individual terms that only depend on the positions of small subsets of the atoms. Moreover, these terms commonly have fixed functional forms describing the various interactions. This simplification is generally called "molecular mechanics" (MM).*

In typical molecular mechanics force fields [35–40], the potential energy is made up of bonded (covalent) and non-bonded terms, $U = U_{\text{bonded}} + U_{\text{non-bonded}}$.

The bonded part usually consists of terms involving the bond lengths b_i between two atoms, the bond angles θ_j between three atoms, and the dihedral angles χ_k between four atoms. A representative functional form is

$$U_{\text{bonded}} = \sum_i^{\text{bonds}} \tfrac{1}{2} K_i^b (b_i - b_{0,i})^2 + \sum_j^{\text{angles}} \tfrac{1}{2} K_j^\theta (\theta_j - \theta_{0,j})^2 + \sum_k^{\text{dihedrals}} K_k^\chi \left[1 + \cos(n_k \chi_k - \chi_{0,k})\right],$$

where K^b, K^θ, K^χ and b_0, θ_0, χ_0, n are the parameters defining the interactions.

The bond and angle terms are commonly described by harmonic "springs". Terms involving dihedral angles can describe both actual dihedral angles formed by three con-

*It is also possible to treat a selected small region quantum-mechanically, and simulate a combined QM/MM model.

secutive bonds as well as so-called "improper" dihedral angles that do not correspond to consecutive bonds but are included to fix the structure, for example, to ensure an aromatic ring remains planar.

The non-bonded part usually consists of electrostatic (Coulomb) and van-der-Waals (dispersion) interactions between pairs of atoms:

$$U_{\text{non-bonded}} = \sum_{\substack{i,j \\ i<j}}^{\text{atoms}} \frac{q_i q_j}{4\pi\epsilon r_{ij}} + \sum_{\substack{i,j \\ i<j}}^{\text{atoms}} \left(\frac{A_{ij}}{r_{ij}^{12}} - \frac{B_{ij}}{r_{ij}^{6}} \right),$$

where r_{ij} is the distance between two atoms, q_i are partial (point) charges, and A_{ij} and B_{ij} are parameters of the Lennard–Jones potential that describes the van-der-Waals interactions. (In recent years there have been extensions to polarisable force fields [41–43] and reactive force fields [44–46], but these are beyond the scope of this thesis.)

The parameters such as spring force constants, Lennard–Jones coefficients, and so forth are determined by a combination of fitting to results from quantum-mechanical calculations, experimental data (*e.g.*, structural data derived from X-ray diffraction), and further *ad-hoc* adjustments to reproduce specific experimental properties [1].

Even though the molecular mechanics potential is a poor approximation of the true quantum-mechanical energy, this approach has been used successfully in vast numbers of simulations of materials and biomolecular systems from small molecules up to proteins, nucleic acids, and lipid membranes [47–49].

Typically, we will be interested in simulations of bulk systems. To get rid of unwanted surface effects, periodic boundary conditions are imposed on the simulation box. The system size needs to be sufficiently large so that the interactions of the periodic mirror images of a particle with themselves do not play a significant role.

The bonded interactions, by their nature, generally involve only a small number of atoms, and their contribution to the total computational cost of force evaluations is typically negligible. In contrast, the evaluation of the non-bonded interactions takes most of the computing time. Whereas the Lennard–Jones potential is usually cut off at around 10 Å without sacrificing accuracy, the Coulomb interaction only decays with the inverse of the distance and cannot be cut off at short range. Algorithms such as the Particle Mesh Ewald method [50] make it feasible to calculate electrostatic interactions in periodic simulation boxes. However, the evaluation of the non-bonded forces is still the bottleneck for many all-atom simulations.

Often, we are interested in the behaviour of a small number of solute molecules within a large volume of solvent. For example, to simulate biomolecules in their natural environment, we need to include a description of the water surrounding them. This can be achieved most accurately by including the ions and water molecules explicitly. In such models, the solvent molecules typically contribute about 90 % of the atoms, and most of the computing resources are spent on calculating the solvent–solvent non-bonded interactions.

However, a dipolar solvent such as water and the dissolved ions will strongly screen even the electrostatic interaction between charged solutes, and the *effective* interactions in bulk simulations generally have a finite, comparatively short range. In physiological conditions, this effective range is around one nanometre, based on the Debye–Hückel screening length for typical ion concentrations [51]. This motivates the use of implicit solvent models in which the solvent contribution to solute–solute interactions is captured in effective interactions. In typical implicit solvent approaches for all-atom models, the charge-screening effects of water are modelled by an infinite continuum medium [52]. However, this does not capture hydrophobic effects properly [53, 54].

2.3. Constant-Temperature Simulations and Thermostats

As discussed previously, evolving a system in time using Eq. (2.1.2) with a conservative force field yields a constant-energy simulation that samples the microcanonical ensemble. This would be appropriate for a system that is completely isolated.

In practice, it would be very difficult to completely isolate a system from its environment, and systems of interest typically exchange heat with their surroundings. We can model this by introducing a heat bath connected to the system. The system exchanges energy with the heat bath, and thereby maintains a constant *temperature* T. This is described by the canonical (NVT) ensemble. To describe a system at constant temperature, the simulation needs to include a thermostat.

A thorough overview and comparison of thermostat algorithms is given by Hünenberger [55]. The two most common approaches are the Nosé–Hoover thermostat and Langevin dynamics.

The Nosé–Hoover thermostat is a fully deterministic approach. It introduces an additional, artificial degree of freedom to the system which represents the heat bath. However, this thermostat has issues related to ergodicity [56, 57], and it can lead to distorted dynamics, especially in small systems.

Langevin dynamics is fully ergodic and samples the canonical ensemble properly, and hence this is the approach we adopt in the present work. It is a stochastic approach that modifies the equations of motion (2.1.1) by adding a friction (damping) term, proportional to the velocities of the atoms, as well as a stochastic noise term (based on the fluctuation–dissipation theorem [58]):

$$m\boldsymbol{a} = \boldsymbol{f} - \gamma\boldsymbol{v} + \sqrt{2\gamma k_\mathrm{B}T}\,\boldsymbol{\eta}(t),$$

where γ is the friction coefficient (typically scaled with the particle mass), k_B is the Boltzmann constant, and $\boldsymbol{\eta}(t)$ describes white noise [with zero mean, $\langle\boldsymbol{\eta}(t)\rangle = 0$, and correlation $\langle\eta_\mu(t)\eta_\nu(t')\rangle = \delta_{\mu\nu}\delta(t-t')$, where $\delta(t-t')$ is the Dirac delta]. In the limit of $\gamma \to 0$, we recover Newtonian MD. The value of γ needs to be chosen large enough for the thermostat to work effectively and small enough not to overdamp low-frequency modes; a reasonable choice is $\gamma = m \times 5\,\mathrm{ps}^{-1}$, where the friction coefficient is proportional to the mass m of each particle [59].

A similar approach which also considers the history of the system and correlations in the noise can be used to obtain "real dynamics" in CG models, where parameters like γ and the noise correlations are derived from underlying atomistic simulations [23].

2.4. Free Energy

When a system is kept at a constant temperature T by letting it exchange energy with a heat bath, its total internal energy E (kinetic + potential energy) is no longer a conserved quantity. Instead, a crucial role is played by the *free* energy A, related to the total energy through $A = E - TS$, where S is the entropy. The dynamics of the system act to minimise the free energy such that, in thermodynamic equilibrium, the free energy takes its minimum value.

Statistical mechanics defines the free energy at a microscopic level as

$$A = -k_\mathrm{B}T\log z,$$

where z is the canonical partition function. For a system of N classical particles with positions $\boldsymbol{r}^N$ and momenta $\boldsymbol{p}^N$, the partition function is given by

$$z = \frac{1}{N!h^{3N}}\int \mathrm{d}^N\boldsymbol{p}\;\mathrm{d}^N\boldsymbol{r}\;\mathrm{e}^{-\beta H(\boldsymbol{r}^N,\boldsymbol{p}^N)},$$

where H is the Hamiltonian and $\beta = 1/k_\mathrm{B}T$. With the Hamiltonian separating into a

sum of kinetic and potential energy, the integral over the momenta can be carried out explicitly. For identical particles, this results in

$$z = \frac{z_N}{N!\Lambda^3},$$

where z_N is the classical configuration integral,

$$z_N = \int \mathrm{d}^N \boldsymbol{r}\ \mathrm{e}^{-\beta U(\boldsymbol{r}^N)}, \tag{2.4.1}$$

and $\Lambda = \left(h^2/(2\pi m k_\mathrm{B} T)\right)^{N/2}$ contains the contribution of the momenta. For a given temperature, this contribution of the momentum integral to the free energy is constant. Therefore, we generally focus on the configurational part, and in the following we will suppress the subscript N and refer to configurational integral and partition function interchangeably.

2.4.1. Collective Variables and the Potential of Mean Force

The integral (2.4.1) cannot be evaluated analytically except for the simplest systems. All practical approaches rely on numerical methods. While there have been attempts to calculate the *absolute* value of the free energy [60], in practice this is rarely feasible, and usually we only care about the free energy differences between states. This can be due to different types of atoms, corresponding to a "transmutation" or alchemical transformation of one molecule into another [61], or due to a transition along a reaction coordinate [62]. For this reason we are often interested in the free energy as a function of parameters that can characterise different states of interest on a larger scale than individual atoms.

These parameters are often referred to as Collective Variables (CVs); they are a function of the atomic configuration, and hence denoted as $\xi(\boldsymbol{r}^N)$. A simple example would be the distance between a selected pair of atoms, $\xi(\boldsymbol{r}^N) = \|\boldsymbol{r}_j - \boldsymbol{r}_i\|$. However, a collective variable does not have to be a geometric parameter; it could be any function of the atomic coordinates, such as a coordination number or the potential energy [63].

For a given collective variable ξ, we can define a "conditional" free energy,

$$A(\xi) = -k_\mathrm{B} T \log z(\xi), \tag{2.4.2}$$

where

$$z(\xi') = \int \mathrm{d}^N \boldsymbol{r}\ \mathrm{e}^{-\beta U(\boldsymbol{r}^N)}\, \delta\big(\xi' - \xi(\boldsymbol{r}^N)\big)$$

is the conditional partition function, corresponding to the configurational integral over all degrees of freedom orthogonal to the collective variable ξ. The probability for the system to be in a state with a certain value of the collective variable ξ is given by

$$p(\xi) = \frac{z(\xi)}{z} = \frac{1}{z} \exp\big[-\beta A(\xi)\big],$$

in which the total partition function $z = \int \mathrm{d}\xi\, z(\xi)$ acts as a normalisation factor. These definitions straightforwardly generalise to a multidimensional set of CVs, $\boldsymbol{\xi}(\boldsymbol{r}^N) = \big\{\xi_i(\boldsymbol{r}^N)\big\}$.

The conditional free energy is commonly referred to as the Potential of Mean Force (PMF): its negative gradient, $-\nabla_\xi A(\xi)$, describes the mean force (as an average over atomic forces) acting on a collective variable. The PMF describes the free energy surface as a function of the CVs, which allows us to gain a better understanding of the stable states of a system and the transition pathways between them. A common use-case for this is a Ramachandran plot [64] of the dihedral backbone angles in peptides which indicates equilibrium configurations and transitions from one to another. For this reason, di- and trialanine are often used in benchmarking of free energy methods [65, 66].

The free energy surface or PMF of a system can be computed by a variety of methods. The most popular methods are the gradient-based Umbrella Integration [67, 68], the Blue Moon method [69, 70], the Adaptive Biasing Force (ABF) method [71, 72], and the histogram-based Metadynamics [73–75]. Gradient-based methods typically estimate the mean forces on a grid and then integrate these results to obtain the PMF. This means that their computational cost scales exponentially with the number of variables; hence, they are usually only applied to up to a few CVs at once. Histogram-based methods directly estimate the probability density function of the CVs (or its logarithm) by accumulating the visited CV values. In Metadynamics, this is based on a history-dependent biasing potential built up of Gaussian peaks that reflects the free energy surface; it can access somewhat larger sets of CVs, though this is typically still limited to no more than six CVs. In a recent study, related to the work in this thesis, ABF and Metadynamics were combined with Gaussian process regression to carry out grid-free multidimensional integration of derivative observations of the free energy [66]. We will review Gaussian process regression in Chapter 4. In the following section we describe how we can obtain mean forces.

2.4.2. Mean Forces

The mean forces (MF) for a given configuration of collective variables can be obtained directly in restrained or constrained MD simulations.

In a restrained MD simulation, a restraint potential U_{rst} for the collective variable is added to the potential energy of the system. This restraint potential is typically harmonic:

$$U_{\mathrm{rst}}(\boldsymbol{r}^N) = \tfrac{1}{2} k_{\tilde{\xi}} \big(\xi(\boldsymbol{r}^N) - \tilde{\xi}_0\big)^2,$$

where $\tilde{\xi}_0$ is the centre of the restraint and $k_{\tilde{\xi}}$ is the restraint force constant. This keeps the value of the CV confined around the position of the restraint. (Confusingly, this was historically referred to as a "harmonic constraint".) We then perform a MD simulation to obtain the thermal average of the CV, $\langle \xi \rangle$, in the presence of the restraint force. This is related to the MF acting on the collective variable ξ by

$$-\frac{\partial A}{\partial \xi}\big(\langle \xi \rangle\big) = k\big(\langle \xi \rangle - \tilde{\xi}_0\big).$$

Note that this calculates the MF at the *average* (biased) value of the CV, and not at the minimum $\tilde{\xi}_0$ of the restraint potential.* As the average value of the CV is shifted proportionally to the magnitude of the mean force, it is difficult to obtain accurate mean forces using restrained MD for unfavourable configurations of CVs. We could reduce $\langle \xi \rangle - \tilde{\xi}_0$ by increasing the strength of the restraint potential, but this would require smaller time steps and increase the error in the measurement.

In a constrained MD simulation, the values of the collective variables are held exactly constant. (This formally corresponds to an "infinite" restraint force.) A constraint solver calculates the Lagrange multipliers that ensure that in each integration step the value of the CVs remains unchanged, resulting in instantaneous "constraint forces". The mean force is given by the ensemble average over the instantaneous constraint forces, represented by a time average over the course of an MD simulation [4]. This is the method on which we rely in this thesis.

Alternatively, it is possible to obtain noisy samples of the free energy gradient by calculating the so-called Instantaneous Collective Forces (ICF) acting on the CVs. We will discuss this in more detail in Section 3.5.

*It is also possible to calculate the mean force at $\tilde{\xi}_0$ [68], but the "unbiasing" from $\langle \xi \rangle$ to $\tilde{\xi}_0$ leads to a much larger error in the value of the mean force [76], and in preliminary investigations we found this unsuitable for our purposes.

2.5. Summary

In this chapter we discussed the principles underlying classical molecular dynamics in the context of all-atom simulations. We introduced the free energy that characterises the state of a system at constant temperature, relevant to typical experimental conditions. Often, we are interested in the properties of a system not as a function of the atomic positions, but as a function of "global" parameters, depending on the atomic positions, that may be more appropriate for describing the configuration on a larger scale. This motivated the introduction of the conditional free energy/potential of mean force as a function of collective variables. In the next chapter, we make the connection between coarse-graining and the free energy by considering the coordinates of the CG model as collective variables in the atomistic model.

3. Coarse-Graining Approaches

As discussed in the introduction, coarse-graining (CG) aims to increase the length and time scales that can be reached in molecular simulations by describing the system at a lower resolution than individual atoms. This reduces the number of interactions that need to be calculated and allows for a larger integration time step, and can thereby speed up the simulation.

The coarse-graining process is necessarily an approximation, and as such it will always require a compromise between various concerns such as speed and use of computational resources, accuracy with respect to different properties (depending on what the modeller is interested in), and transferability between different systems and state points.

The construction of a CG model consists of two tasks. First, we need to define the CG sites and their topology*. Second, we need to specify the interactions between CG sites. The definition of CG sites depends closely on the system of interest and the questions that the model should answer. There is a wide range of options for how to define the sites; a CG site can describe a single atom, a "united atom" that describes a heavy atom together with all the hydrogen atoms bound to it, a small group of atoms, a single side chain, a residue or whole molecule, or even several molecules at once. The choice of CG mapping influences the properties that can be reproduced by the CG simulation, and hence this forms an important part of the modelling process. So far, there has been only limited research into a systematic determination of CG sites [77–79], and the definition of CG sites is predominantly based on physical/chemical intuition. The definition of CG sites depends on the level of description we seek, and should consider several factors [80, 81]: First, the phenomena in which we are interested need to be representable by the coarse-grained description, and the CG sites should be able to capture the slow, large-amplitude motions in the system. Second, to speed up the simulations, the choice of CG sites should eliminate unnecessary detail and filter out the fast, high-frequency but low-amplitude motions. Moreover, the choice of CG

*The topology describes the structure of the molecules; it defines the bonded interactions between sites. Two CG sites are generally considered "bonded" when there are covalent bonds connecting atoms from both sites.

sites impacts how well the effective coarse-grained interactions are able to reproduce properties of the atomistic system. There should be a one-to-one mapping between distinct molecular states in the all-atom model and in the CG description [82], and coarse-graining will produce more accurate results when the CG degrees of freedom are less correlated [79]. (Hence, increasing the resolution of the CG mapping may actually lead to reduced accuracy [22,83].)

In this thesis, we will assume the definition of CG sites as given. We focus on the second task in creating a CG model, that is, specifying the interactions between CG sites. Whereas in all-atom modelling, empirical force fields are now being routinely applied in a production setting, for each new CG model the CG interactions need to be parametrised again (as they depend on the choice of sites and on the state point). How to generate CG interactions is still a topic of active research.

There are many different approaches to coarse-graining. They can generally be classified into more phenomenological "top-down" models and "bottom-up" models that are based on an underlying all-atom model. This distinction will be discussed in Section 3.1. Bottom-up models rely on a mapping from the all-atom to the coarse-grained description; we will introduce this mapping in Section 3.2. The mapping between all-atom and CG model allows us to consider the *consistency* of the CG model in terms of probability distributions. In Section 3.3 we will discuss this and see how the Potential of Mean Force (PMF) for the coarse-grained coordinates (*cf.* Section 2.4.1) provides a CG potential that exactly reproduces all structural and thermodynamic properties of the underlying all-atom model [79,84]. However, using the many-body PMF as the interaction potential for the CG model is generally not feasible.

In practice, coarse-graining approaches rely on simpler potentials that describe effective interactions, optimised to reproduce a subset of the properties of the underlying system. In the second half of this chapter we review and compare the current approaches to determining coarse-grained interactions that can be found in the literature. We discuss structure-based approaches in Section 3.4 and force-based approaches in Section 3.5. We conclude with a summary in Section 3.6.

3.1. Top-down *vs.* Bottom-up Models

CG models are commonly classified into whether they follow a "top-down" or a "bottom-up" approach. Here we outline and compare these two coarse-graining philosophies. A more extensive discussion (including other, more specialised approaches such as "knowledge-based" network models) can be found, for example, in Ref. 80.

Top-down models

In the top-down approach, interaction terms between CG sites are proposed based on physical principles (with no connection to an underlying atomistic model). Typically, the interactions are given by fixed, analytical functional forms. Their parameters can be optimised to reproduce quantities that are measured in experiments, such as thermodynamic properties, or are simply specified based on physical intuition. Top-down models can be considered "fine-grained" models, parametrised on (experimental) data from larger scales.

Top-down models generally focus on reproducing thermodynamic and other large-scale properties. They are often seen as a "minimal model" to elucidate what types of interactions need to be considered to reproduce specific properties and emergent behaviours in a system. (However, different underlying atomistic interactions can lead to the same patterns on a larger scale, so top-down models cannot explain how or which atomistic properties lead to these patterns.)

Top-down models have been successfully used in practice for a wide range of systems [17]. They are well suited for studies of large changes, such as DNA duplex formation [85,86]. Moreover, they tend to be comparatively transferable between state points [17]. However, top-down models are often not quantitatively accurate. Generally, they cannot be improved systematically, and their relation to more detailed models and properties at the atomistic scale is undetermined.

Bottom-up models

Alternatively, we can base a CG model on an underlying, more detailed "fine-grained" model, with the aim of inferring the CG interactions from simulation results of the more detailed model; this is called the bottom-up approach. This takes advantage of all the effort that went into building and validating the fine-grained model and removes many of the arbitrary choices that are made in the development of top-down models. It is common, but not necessary, that the underlying fine-grained model is an all-atom model. In principle, this allows a hierarchical approach to modelling, starting on the atomistic scale and going up to a very large-scale CG model. In the following, we assume that the fine-grained model is an all-atom model.

The defining feature of bottom-up coarse-graining is that there is a *mapping* function from the all-atom model to the CG model. This is a many-to-one mapping: multiple all-atom configurations correspond to the same coarse-grained configuration. We consider this mapping function in more detail in Section 3.2. The explicit mapping

makes it possible to systematically connect models at different scales ("multiscale coarse-graining").

A disadvantage of the bottom-up approach is that the CG model can only contain information that has been sampled in the atomistic model: If the all-atom trajectory that is used to construct the CG model does not explore the full phase space but is caught in a free-energy basin, the CG model cannot represent the behaviour of the atomistic model outside this basin (irrespective of the approach used to determine the CG interactions). Generally, this means the CG model will only reproduce the behaviour inside this basin, even if, on *long* time scales, there would be fluctuations that take the all-atom system out of this basin. Hence, bottom-up models are often biased towards the structures for which they were parametrised. In principle, they are primarily suited to describe fluctuations around the equilibrium structure. (We note that there are free energy methods such as Metadynamics [75] and ABF [87] with which we can run *biased* simulations that enable the model to escape local free energy basins and explore wider parts of the free energy surface, though this is beyond the scope of this thesis.)

Comparison

Noid [80] characterises the distinction between these types of approaches as follows: top-down models are *underconstrained* by experimental data — there is a large degree of flexibility in terms of what functional forms and parameters to choose, and there are not enough different experimental observables to uniquely determine them —, whereas bottom-up models are *overconstrained* by all the data available from simulations of the underlying all-atom model. Bottom-up models tend to focus on structural properties, whereas top-down models tend to focus on large-scale emergent phenomena.

Top-down modelling has generally been used for low-resolution models (with each CG site describing a large part of a residue or even several residues at once), whereas bottom-up models have historically been used only for "soft", limited coarse-graining (with a few atoms per CG site). However, recently there has been more interest in bottom-up approaches to drastic coarse-graining of molecules like lipids where each CG site encompasses thirty to fifty atoms [88]. This is where we expect the GAP-CG approach that we introduce in this thesis to be most useful.

We are interested in a consistent bridge between scales, following on from previous research to connect atomistic molecular dynamics with quantum-mechanical potentials [29]. A link between CG simulations and the underlying atomistic structure is crucial,

for example, for developing new materials with specific properties [89]. For this reason, we will focus on bottom-up approaches to coarse-graining.

In the following, we first discuss the mapping between the underlying all-atom model and the coarse-grained description (Section 3.2). We then consider a definition of *consistency* between all-atom and CG model (Section 3.3). In the remainder of this chapter, we discuss different approaches to determining CG interactions from simulations of the underlying all-atom model, based on its structural distributions (Section 3.4) or on the atomistic forces (Section 3.5).

3.2. Mapping from All-Atom to Coarse-Grained Model

The defining feature of bottom-up CG models is their explicit connection to an underlying (atomistic) model, given by a mapping from the all-atom model to the CG model. Here we show how this connection is formally made.

In the remainder of this chapter, variables denoting the atomistic model are designated by lower-case characters, and variables denoting the CG model are designated by upper-case characters.

The state of the atomistic system is specified by the set of positions and momenta of all n atoms. We collectively describe this state by the $3n$-dimensional vectors for position, $\boldsymbol{r}^n$, and momenta, $\boldsymbol{p}^n$:

$$\boldsymbol{r}^n = \{\boldsymbol{r}_1, \ldots, \boldsymbol{r}_n\}, \qquad \boldsymbol{p}^n = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\}.$$

Correspondingly, the state of the CG model is defined by the positions and momenta of the N CG sites, where we assume that the CG sites are defined as structureless, point-like particles:

$$\boldsymbol{R}^N = \{\boldsymbol{R}_1, \ldots, \boldsymbol{R}_N\}, \qquad \boldsymbol{P}^N = \{\boldsymbol{P}_1, \ldots, \boldsymbol{P}_N\}.$$

We now introduce the mapping function from the atomistic configuration $\boldsymbol{r}^n$ to the CG configuration $\boldsymbol{R}^N$:

$$\boldsymbol{R}^N = \boldsymbol{M}_{\boldsymbol{R}}^N(\boldsymbol{r}^n), \tag{3.2.1}$$

where $\boldsymbol{M}_{\boldsymbol{R}}^N = \{\boldsymbol{M}_{\boldsymbol{R}1}, \ldots, \boldsymbol{M}_{\boldsymbol{R}N}\}$. Note that this is a many-to-one mapping: multiple atomistic configurations correspond to the same CG configuration.

In principle, Eq. (3.2.1) could describe a non-linear mapping, but in the following we only consider linear mappings, as is prevalent in the literature. This simplifies many

equations and makes it easier to reason about and compare the different approaches to coarse-graining. A linear mapping is given by

$$\boldsymbol{M}_{\boldsymbol{R}I}(\boldsymbol{r}^n) = \sum_{i=1}^{n} c_{Ii}\boldsymbol{r}_i \qquad \text{for } I = 1, \ldots, N. \tag{3.2.2}$$

We require translational invariance: if we shift all atomistic coordinates by a certain amount, the CG coordinates should shift by the same amount. This imposes the constraints $\sum_{i=1}^{n} c_{Ii} = 1$ for all I. It will be useful to introduce the set $\mathcal{I}_I$ of all atoms that are *involved* in the definition of a CG site I, that is, all atoms i for which $c_{Ii} \neq 0$. Similarly, the set $\mathcal{S}_I$ contains all atoms *specific* to a site I, that is, all atoms which are not involved in any other CG site. Common choices for the definition of CG sites are the centre of mass ($c_{Ii} = m_i / \sum_{j\in\mathcal{I}_I} m_j$) or the centre of geometry ($c_{Ii} = 1 / \sum_{j\in\mathcal{I}_I} 1$) of a group of atoms.

Corresponding to the mapping function (3.2.1) for the coordinates, there is a similar mapping from the momenta of the atoms to the momenta of the CG sites,

$$\boldsymbol{P}^N = \boldsymbol{M}_{\boldsymbol{P}}^N(\boldsymbol{p}^n),$$

where $\boldsymbol{M}_{\boldsymbol{P}}^N = \{\boldsymbol{M}_{\boldsymbol{P}1}, \ldots, \boldsymbol{M}_{\boldsymbol{P}N}\}$. Coarse-grained position and momentum are connected by $\boldsymbol{P}_I = M_I \, \mathrm{d}\boldsymbol{R}_I / \mathrm{d}t$, where M_I is the mass of CG site I. Hence, the linear mapping (3.2.2) induces the following relation for the CG momenta:

$$\boldsymbol{M}_{\boldsymbol{P}I}(\boldsymbol{p}^n) = M_I \sum_{i=1}^{n} c_{Ii}\boldsymbol{p}_i / m_i \qquad \text{for } I = 1, \ldots, N.$$

For a typical CG model, $N \ll n$, and thus it has far fewer degrees of freedom than the corresponding all-atom model. We call the $3(N-n)$ degrees of freedom that are lost in the coarse-graining the "internal" degrees of freedom. The reduction in the number of degrees of freedom brings with it an entropy difference between CG and all-atom model [84]. In the following section, we consider consistency between all-atom and coarse-grained description, and the conditions this imposes on the CG model.

3.3. Consistency and the Many-Body Potential of Mean Force

Having defined a mapping function, we can now analyse the relation between the atomistic and the coarse-grained scales within a statistical mechanics framework. Here we consider the phase space probability distribution of both models, following Ref. 90. A more detailed discussion can be found in Ref. 91.

The phase space probability distribution of the atomistic system is determined by its Hamiltonian H_{AA}; we assume that it can be written as the sum of momentum-dependent kinetic energy and position-dependent potential energy:

$$H_{\mathrm{AA}} = \sum_{i=1}^{n} \frac{\boldsymbol{p}_i^2}{2m_i} + u(\boldsymbol{r}^n).$$

Then the phase space probability distribution factorises into independent contributions from coordinates and momenta,

$$p_{rp}(\boldsymbol{r}^n, \boldsymbol{p}^n) = p_r(\boldsymbol{r}^n) p_p(\boldsymbol{p}^n),$$

where

$$p_r(\boldsymbol{r}^n) \propto \exp\bigl[-u(\boldsymbol{r}^n)/k_{\mathrm{B}}T\bigr] \qquad \text{and} \tag{3.3.1a}$$

$$p_p(\boldsymbol{p}^n) \propto \exp\bigl(-\sum_{i=1}^{n} \boldsymbol{p}_i^2/(2m_i k_{\mathrm{B}}T)\bigr) \tag{3.3.1b}$$

are the probability distributions of coordinates and momenta, respectively (defined up to a normalisation factor).

Analogously, the CG model is described by the Hamiltonian

$$H_{\mathrm{CG}} = \sum_{I=1}^{N} \frac{\boldsymbol{P}_I^2}{2M_I} + U(\boldsymbol{R}^N),$$

and the probability distribution of the CG model factorises as

$$P_{RP}(\boldsymbol{R}^N, \boldsymbol{P}^N) = P_R(\boldsymbol{R}^N) P_P(\boldsymbol{P}^N),$$

where

$$P_R(\boldsymbol{R}^N) \propto \exp\big[-U(\boldsymbol{R}^N)/k_\mathrm{B}T\big] \qquad \text{and} \tag{3.3.2a}$$

$$P_P(\boldsymbol{P}^N) \propto \exp\big(-\sum_{I=1}^{N} \boldsymbol{P}_I^2/(2M_I k_\mathrm{B}T)\big). \tag{3.3.2b}$$

We now make use of the mapping function and consider the probability of finding a given CG configuration within the atomistic ensemble, $p_R(\boldsymbol{R}^N)$. This is obtained by integrating out the internal degrees of freedom, averaging over their fluctuations:

$$p_R(\boldsymbol{R}^N) = \int \mathrm{d}\boldsymbol{r}^n \, p_r(\boldsymbol{r}^n)\, \delta\big(\boldsymbol{M}_{\boldsymbol{R}}^N(\boldsymbol{r}^n) - \boldsymbol{R}^N\big), \tag{3.3.3}$$

where

$$\delta\big(\boldsymbol{M}_{\boldsymbol{R}}^N(\boldsymbol{r}^n) - \boldsymbol{R}^N\big) \equiv \prod_{I=1}^{N} \delta\big(\boldsymbol{M}_{\boldsymbol{R}I}(\boldsymbol{r}^n) - \boldsymbol{R}_I\big)$$

is the $3N$-dimensional Dirac delta function, ensuring that we only include those atomistic configurations that map to the given CG configuration.

Similarly, for the CG momenta we obtain

$$\begin{aligned} p_P(\boldsymbol{P}^N) &= \int \mathrm{d}\boldsymbol{p}^n \, p_p(\boldsymbol{p}^n)\, \delta\big(\boldsymbol{M}_{\boldsymbol{P}}^N(\boldsymbol{p}^n) - \boldsymbol{P}^N\big) \\ &\propto \int \mathrm{d}\boldsymbol{p}^n \, \Big(\prod_i \exp\big(-\boldsymbol{p}_i^2/(2m_i k_\mathrm{B}T)\big)\Big) \prod_I \delta\big(M_I \sum_j c_{Ij}\boldsymbol{p}_j/m_j - \boldsymbol{P}_I\big). \end{aligned} \tag{3.3.4}$$

The configuration and momentum terms are independent of each other and can be treated separately. We discuss the momentum term in Section 3.3.1 and the configuration term in Section 3.3.2.

3.3.1. Momentum Space

We consider a CG model to be consistent with an underlying atomistic model if the probability distribution in the CG model matches the probability distribution of the CG sites in the atomistic model, as obtained through the corresponding mapping function. That is, for consistency in momentum space we require (3.3.4) and (3.3.2b) to be equal:

$$P_P(\boldsymbol{P}^N) = p_P(\boldsymbol{P}^N). \tag{3.3.5}$$

The left-hand side of this equation can be written as a product of independent zero-mean Gaussians for each I,

$$P_P(\boldsymbol{P}^N) \propto \prod_I \exp\bigl(-\boldsymbol{P}_I^2/(2M_I k_\mathrm{B}T)\bigr).$$

Requiring this to be equivalent to the probability of the CG momenta within the atomistic model imposes two conditions on the CG model: First, for $p_P(\boldsymbol{P}^N)$ to be a product of independent zero-mean Gaussians as well, no atom can be involved in the definition of more than one CG site. Second, requiring the second moments of the Gaussians in both probability distributions to be equal places the following condition on the CG masses M_I [90]:

$$\frac{1}{M_I} = \sum_i \frac{c_{Ii}^2}{m_i} \qquad \text{for all } I.$$

Of all possible definitions of CG sites for a given group of atoms, identifying the site with the centre of mass of the atoms results in the largest mass for the CG site. As this corresponds to slower motions this will often be the preferred choice.

Note that consistency in momentum space only implies that the momentum distribution in the CG model matches the one in the atomistic model. It does not imply that the CG model will produce the correct dynamics (for example, time correlation functions). Obtaining "real dynamics" [23] is beyond the scope of this work, and in the remainder of this thesis we focus on consistency in configuration space.

3.3.2. Configuration Space

Analogously to Eq. (3.3.5), we consider the CG model to be consistent with the underlying atomistic model in configuration space if the probability distributions of the CG coordinates in the CG model (3.3.2a) and in the atomistic ensemble (3.3.3) are equal. That is, for consistency we require

$$P_R(\boldsymbol{R}^N) = p_R(\boldsymbol{R}^N). \tag{3.3.6}$$

This is equivalent to

$$\exp\bigl[-U(\boldsymbol{R}^N)/k_\mathrm{B}T\bigr] \propto \int \mathrm{d}\boldsymbol{r}^n \, \exp\bigl[-u(\boldsymbol{r}^n)/k_\mathrm{B}T\bigr] \, \delta\bigl(\boldsymbol{M}_{\boldsymbol{R}}^N(\boldsymbol{r}^n) - \boldsymbol{R}^N\bigr),$$

thereby defining the "optimal" CG potential $U(\boldsymbol{R}^N)$. Note that this is not a potential energy, but also contains an entropic contribution due to integrating over fluctuations in the internal degrees of freedom. Hence, this CG potential actually describes a free energy surface, and it explicitly depends on the temperature T. We can explicitly write U as a conditional free energy:

$$U(\boldsymbol{R}^N) = -k_\mathrm{B}T \log z(\boldsymbol{R}^N) + (\text{const.}), \tag{3.3.7}$$

where

$$z(\boldsymbol{R}^N) \equiv \int \mathrm{d}\boldsymbol{r}^n \, \exp\big[-u(\boldsymbol{r}^n)/k_\mathrm{B}T\big] \, \delta\big(\boldsymbol{M}_{\boldsymbol{R}}^N(\boldsymbol{r}^n) - \boldsymbol{R}^N\big) \tag{3.3.8}$$

is the constrained partition function (as a function of the CG coordinates). Hence, Eq. (3.3.7) is equivalent to the Potential of Mean Force (PMF) $A(\xi)$ discussed in the previous chapter [Eq. (2.4.2) in Section 2.4, page 21]. Whereas the PMF is usually considered for a small number of collective variables ξ, such as order parameters, reaction coordinates or pair distances, the CG potential corresponds to the *many-body* PMF $A(\boldsymbol{R}^N)$ for all $3N$ CG coordinates.

When the CG interactions are given by the PMF for the CG coordinates, both structural and thermodynamic properties of the atomistic model are exactly preserved. However, explicitly evaluating the integral in Eq. (3.3.8) would not be feasible in practice. Moreover, $A(\boldsymbol{R}^N)$ describes many-body interactions that, in principle, involve all CG sites simultaneously. In the following, we review current approaches for obtaining CG potentials that approximate Eq. (3.3.7) or otherwise generate effective interactions that aim to reproduce probability distributions of the atomistic model.

3.4. Structure-based Approaches

Structure-based approaches to coarse-graining derive the CG interactions from structural distributions of the underlying system. We give an overview of the four main approaches that can be found in the literature: Direct and Iterative Boltzmann Inversion (DBI and IBI; Sections 3.4.1 and 3.4.2, respectively), Inverse Monte Carlo (IMC; Section 3.4.3), and Relative Entropy (Section 3.4.4).

Direct and Iterative Boltzmann Inversion are based solely on one-dimensional distributions of individual collective variables (CVs). Inverse Monte Carlo also takes into account cross-correlations between distributions of different CVs. The Relative Entropy approach considers, in principle, the joint probability distribution of all CG coordinates. In any case, typically only pair potentials are considered as the basis for

the (non-bonded) CG interactions. If there are only pairwise interactions, the Henderson theorem [92] shows that the pair potential that reproduces the radial distribution function (RDF) is unique (up to a constant). However, it should be noted that in practice a series of different potentials may result in coinciding RDFs [93, 94].

3.4.1. Direct Boltzmann Inversion

When a collective variable ξ is completely uncorrelated with all other degrees of freedom of the system, it appears as an independent factor in the probability distribution of the CG model [Eq. (3.3.2a)]. In this case, the probability distribution for this CV can be written as a Boltzmann weight,

$$p(\xi) \propto \exp\big(-\beta A_\xi(\xi)\big), \tag{3.4.1}$$

where $A_\xi(\xi) = -k_\mathrm{B}T \log p(\xi) + (\text{const.})$ is the conditional free energy for ξ, and $\beta = 1/k_\mathrm{B}T$. Based on this, Tschöp *et al.* [95] proposed a CG potential for the variable ξ obtained by "Boltzmann inversion":

$$U_\xi(\xi) = -k_\mathrm{B}T \log \frac{p(\xi)}{J(\xi)}. \tag{3.4.2}$$

Note the Jacobian factor $J(\xi)$ that, for the general case when ξ is not a Cartesian coordinate, accounts for the change of volume elements in the integration over orthogonal degrees of freedom. Commonly used collective variables are the distance r between a pair of (bonded or non-bonded) CG sites, the angle θ between three CG sites, and the dihedral angle χ defined by four CG sites. The corresponding Jacobian factors are $J(r) = r^2$, $J(\theta) = \sin(\theta)$ and $J(\chi) = 1$, respectively. In practice, this approach is used even when a CV does have some correlation with other degrees of freedom. However, when multiple ξ_i are correlated, we cannot just add up the corresponding DBI potentials without introducing a bias.

We illustrate this for the example of the pairwise distance $r = \|\boldsymbol{R}_J - \boldsymbol{R}_I\|$ in a one-site model. In this case, the scaled distribution $p(r)/J(r) = p(r)/r^2$ is proportional to the two-body pair correlation function $g^{(2)}(\boldsymbol{R}_I, \boldsymbol{R}_J) = g(r)$, more commonly known as the Radial Distribution Function or RDF. It is related to the *two-body* potential of mean force $w^{(2)}(r)$, sometimes referred to as the pair potential of mean force (ppmf) [96]:

$$g(r) = \exp\big(-\beta w^{(2)}(r)\big).$$

Direct Boltzmann Inversion results in the CG pair potential

$$U_{\mathrm{CG}}(r) = -k_{\mathrm{B}}T \log g_{\mathrm{AA}}(r) = w^{(2)}_{\mathrm{AA}}(r), \tag{3.4.3}$$

where for clarity we added the subscript "AA" to refer to the all-atom model. In the limit of infinite dilution (when there are no correlations between different pairs of particles), a CG simulation using the potential (3.4.3) will exactly reproduce the all-atom RDF, as in this case $w^{(2)}_{\mathrm{CG}} = U_{\mathrm{CG}} = w^{(2)}_{\mathrm{AA}}$. However, in general, the ppmf is a sum of the direct interaction potential and environment-mediated interactions. The two-body pair correlation function contains effective contributions due to many-body effects, such as the packing of non-overlapping particles in the liquid phase.

This can easily be understood by considering a CG model that is based on a one-to-one mapping with a single-site hard-sphere liquid. In the fine-grained model, the interactions are given by a pair potential that is zero beyond the hard sphere radius. Because of the one-to-one mapping, the CG potential should be identical to the atomistic potential. However, the close packing of the spheres leads to oscillations in the RDF, and hence the CG potential obtained by DBI will also show an oscillatory behaviour. This is entirely due to the environmental effect of correlations with other sites. (In contrast, for an ideal gas, the RDF is exactly one beyond the hard sphere radius, correctly resulting in the DBI potential being zero in that region.)

When CVs are significantly correlated, the potentials obtained by DBI will not reproduce the distributions in the atomistic model, and hence cannot be used directly. The assumption of independent degrees of freedom is a good approximation for non-bonded pair potentials when CG sites are *dilute* and for bond potentials when the CG bonds are *stiff* [80]. In practice, DBI is commonly used for bonded interactions (bonds, angles and dihedral angles), and the remaining non-bonded interactions can be obtained by other, more refined methods. Non-bonded interactions can only be successfully described by DBI when the CG sites are very dilute. However, the DBI pair potential can provide a useful starting point for other approaches, such as IBI and IMC, that iteratively refine CG pair potentials to improve the reproduction of the atomistic pair correlation function.

3.4.2. Iterative Boltzmann Inversion

As discussed in the previous section, Direct Boltzmann Inversion (DBI) neglects environmental effects. Hence, in general, it does not reproduce the atomistic distributions in the CG model. Reith *et al.* [97] proposed *Iterative* Boltzmann Inversion (IBI) as an

extension of DBI to derive *effective* pair potentials that do reproduce atomistic distributions. While this approach also applies to other degrees of freedom, we focus on non-bonded pair potentials and the RDF. Starting from an initial guess $U_{(0)}(r)$ for the CG pair potential (for example, based on DBI of the RDF), in each iteration step a correction is applied to the CG potential:

$$U_{(i+1)}(r) = U_{(i)}(r) + \alpha \Delta U_{(i)}(r),$$

where α is an optional scaling factor ($0 < \alpha < 1$) to ensure the stability of the iterative procedure. The correction is given by the difference between effective CG ppmf and all-atom ppmf,

$$\Delta U_{(i)}(r) = w^{(2)}_{\mathrm{CG},(i)}(r) - w^{(2)}_{\mathrm{AA}}(r) = k_{\mathrm{B}}T \log \frac{g_{\mathrm{CG},(i)}(r)}{g_{\mathrm{AA}}(r)},$$

where g_{AA} is the reference RDF of the atomistic model. This process is repeated until the RDF in the CG model is sufficiently close to the all-atom RDF.

This scheme is motivated by the following observation: when $g_{\mathrm{CG},(i)}(r) > g_{\mathrm{AA}}(r)$ at a given pair distance r, this implies that the ppmf of the CG model at that distance is too attractive, and that it should be more repulsive. As the ppmf is given by the sum of direct pair interaction $U_{(i)}(r)$ and environmental contributions, the simplest way to increase the ppmf is to increase the value of the pair potential at that distance. Clearly, $g^{\mathrm{CG}}_{*}(r) = g_{\mathrm{AA}}(r)$ is a fixed point of this iteration. Hence, assuming convergence, the IBI procedure results in a pair potential that reproduces the RDF of the all-atom system.

If more than one type of interactions should be derived in this way, it is advantageous to update one potential at a time while keeping the others fixed, as this leads to improved convergence. This procedure works best by starting out with the stiffest, least correlated degrees of freedom; thus, the typical order would be bond, angle, dihedral, and finally non-bonded interactions [97]. When there are significant cross-correlations between different pair correlation functions, as is the case for the non-bonded degrees of freedom in a multicomponent systems, the IBI procedure may have convergence problems [94]. IBI is a local algorithm in the sense that the update to the pair potential at a distance r only depends on the values of CG and atomistic RDF at that distance, and the potential updates for different interactions are independent of each other. This implies that uncertainty in the RDF at a certain distance only has a constant effect on the potential update, independent of the distance. However, in dilute systems it can be difficult to sample the RDF between sites to sufficient accuracy [98].

3.4.3. Inverse Monte Carlo

To explicitly account for the correlations due to different types of interactions in multicomponent systems, Lyubartsev and Laaksonen proposed the Inverse Monte Carlo (IMC) approach, also referred to as Reverse Monte Carlo (RMC) [99–101]. Like IBI, it is an iterative approach. IMC is based on linear response theory, estimating how a change in the CG pair potential $U_{\tilde{\xi}'}(r')$ for one type of interaction $\tilde{\xi}'$ at distance r' impacts *all* pairwise distance distributions $P_{\tilde{\xi}}(r)$. This is given by the susceptibility matrix,

$$K_{\tilde{\xi}\tilde{\xi}'}(r, r'|U) = \frac{\delta P_{\tilde{\xi}}(r|U)}{\delta U_{\tilde{\xi}'}(r')},$$

defined as the functional derivative of the probability distribution $P_{\tilde{\xi}}(r)$ with respect to the pair potential $U_{\tilde{\xi}'}(r')$. The susceptibility matrix can be related to the covariance between different pair distances,

$$K_{\tilde{\xi}\tilde{\xi}'}(r, r') = \beta\big(P_{\tilde{\xi}}(r)P_{\tilde{\xi}'}(r') - P_{\tilde{\xi}\tilde{\xi}'}(r, r')\big),$$

where $P_{\tilde{\xi}\tilde{\xi}'}(r, r')$ is the *joint* probability distribution (cross-correlation) of finding a pair of type $\tilde{\xi}$ at distance r and finding a pair of type $\tilde{\xi}'$ at distance r'. Corrections to the pair potentials induce changes in the pair distance distributions. This is given by

$$\delta P_{\tilde{\xi}}(r) = \sum_{\tilde{\xi}'} \int \mathrm{d}r'\, K_{\tilde{\xi}\tilde{\xi}'}(r, r')\delta U_{\tilde{\xi}'}(r'), \tag{3.4.4}$$

where $\delta P_{\tilde{\xi}}(r) = p_{\tilde{\xi}}(r) - P_{\tilde{\xi}}(r)$ is the difference between the atomistic and CG distribution. In practice, the pair potentials are discretised, and Eq. (3.4.4) leads to a system of coupled linear equations, whose solution provides the corrections to the pair potentials. For this we need to invert the susceptibility matrix K, which can only be determined approximately. Inverting it could amplify the errors, leading to convergence issues. To alleviate such issues, we can scale the potential correction by a factor $0 < \alpha < 1$ as in IBI, though this may increase the number of iterations required for convergence. This process is repeated until the atomistic distributions are reproduced sufficiently accurately. As a covariance matrix, K is positive definite. Hence, in principle, the IMC algorithm will converge to a global optimum.

Inverse Monte Carlo has a single, global update step for all interactions in the system. As this takes into account cross-correlations between different interactions, this is a particular advantage in multi-component systems, where IMC provides more stable updates than IBI. Furthermore, IMC converges faster than IBI with the number of

iterations [94].

However, whereas IBI only depends on averages of pair distance distributions (the radial distribution functions), IMC also requires us to accurately calculate the cross-correlations between all combinations of pair distances; this necessitates much longer molecular dynamics runs in each iteration step. Moreover, Ref. 94 also demonstrated that IMC suffers from significant finite-size effects. In contrast to IBI, the influence of errors in the distributions on the IMC updates scales quadratically with the distance; for this reason, the long-range tails of the RDFs must be converged to a very high accuracy with respect to both system size and simulation length. Hence, IMC may not be appropriate for small systems.

In principle, IMC could be applied to three-body potentials, but this comes with a prohibitive computational cost [26, 99].

3.4.4. Relative Entropy

Another, recent approach to structure-based coarse-graining was proposed by Shell, based on the relative entropy S_{rel} of a CG model [102]. The relative entropy measures how well we can estimate the underlying atomistic distribution when sampling from a given CG distribution. It is defined as

$$S_{\mathrm{rel}} = \int \mathrm{d}\boldsymbol{r}^n \, p_r(\boldsymbol{r}^n) \log\left(\frac{p_r(\boldsymbol{r}^n)}{P_R\big(\boldsymbol{M}_{\boldsymbol{R}}^N(\boldsymbol{r}^n)\big)} \right) + \langle S_{\mathrm{map}} \rangle_r, \tag{3.4.5}$$

where

$$S_{\mathrm{map}}(\boldsymbol{R}^N) = \log \int \mathrm{d}\boldsymbol{r}^n \, \delta\big(\boldsymbol{M}_{\boldsymbol{R}}^N(\boldsymbol{r}^n) - \boldsymbol{R}^N\big)$$

is called the mapping entropy. S_{map} describes the volume of atomistic configuration space corresponding to a single CG configuration; it depends only on the CG mapping $\boldsymbol{M}_{\boldsymbol{R}}^N(\boldsymbol{r}^n)$, not on the CG force field.

The relative entropy is a measure of the overlap between the atomistic and the CG ensembles, related to the Kullback–Leibler divergence between p_r and P_R; as such, it implicitly depends on the CG potential U_{CG}. It can be considered to describe the amount of information that is lost when moving from an all-atom description to the CG model.

According to the Gibbs inequality,

$$S_{\mathrm{rel}}[U_{\mathrm{CG}}] \geq 0,$$

and the relative entropy only vanishes when the CG and all-atom probability distributions match exactly.

This motivates an optimisation of force field parameters based on minimising the relative entropy. Substituting p_r and P_R [using Eqs. (3.3.1a) and (3.3.2a)], we obtain

$$S_{\mathrm{rel}} = \beta\big\langle U_{\mathrm{CG}} - U_{\mathrm{AA}}\big\rangle_{\mathrm{AA}} - \beta(A_{\mathrm{CG}} - A_{\mathrm{AA}}) + \big\langle S_{\mathrm{map}}\big\rangle_{\mathrm{AA}}, \tag{3.4.6}$$

where the subscripts "AA" and "CG" refer to the all-atom and coarse-grained models, respectively; U is the potential energy and A is the free energy. The angular brackets denote the ensemble average.

The relative entropy can be considered to be a function of all free parameters λ of the CG potential; we can minimise it numerically without requiring the absolute value, guided by the first derivatives [103]:

$$\frac{\partial S_{\mathrm{rel}}}{\partial \lambda} = \beta \left\langle \frac{\partial U_{\mathrm{CG}}}{\partial \lambda} \right\rangle_{\mathrm{AA}} - \beta \left\langle \frac{\partial U_{\mathrm{CG}}}{\partial \lambda} \right\rangle_{\mathrm{CG}}.$$

For parameters that enter the CG potential linearly, the curvature implied by the second derivatives of the relative entropy is always positive, and hence there will only be a single minimum. However, the relative entropy can also be minimised for analytic functional forms and parameters that enter non-linearly.

When obtaining CG interactions by minimising the relative entropy, the resulting interactions can reproduce the expectation values of all terms that are included in the potential. For instance, this means that finely tabulated pair potentials will reproduce the RDF, and in this case relative entropy optimisation is equivalent to IBI or IMC. Moreover, for a CG potential that can accurately represent an N-body potential, minimising the relative entropy will reproduce the PMF [104].

3.5. Force-based Approaches

The approaches reviewed in the previous section focus on the reproduction of the *structure* of the atomistic model to determine effective CG interactions. Alternatively, we can aim to reproduce the mean *forces* acting on the CG sites. This approach is based on the force-matching scheme of Ercolessi and Adams [105] and has been extensively developed by Voth *et al.* [90, 106–108]. Reproducing the mean forces forms the basis for our own approach to coarse-graining (which will be introduced in Chapter 4), hence we will discuss the derivation given in Ref. 90 in more detail.

3.5.1. Instantaneous Collective Forces

We start from the "ideal" CG potential based on the many-body PMF as defined by Eq. (3.3.7). The coarse-grained force on site I is given by the negative gradient of the PMF,

$$\boldsymbol{F}_I(\boldsymbol{R}^N) = -\frac{\partial U(\boldsymbol{R}^N)}{\partial \boldsymbol{R}_I}.$$

Carrying out the derivative, we obtain

$$\boldsymbol{F}_I(\boldsymbol{R}^N) = \frac{k_\mathrm{B}T}{z(\boldsymbol{R}^N)} \int \mathrm{d}\boldsymbol{r}^n \, \exp\big[-u(\boldsymbol{r}^n)/k_\mathrm{B}T\big] \prod_{J(\neq I)} \delta\big(\boldsymbol{M}_{\boldsymbol{R}J}(\boldsymbol{r}^n) - \boldsymbol{R}_J\big) \, \frac{\partial}{\partial \boldsymbol{R}_I} \delta\big(\sum_{i\in\mathcal{I}_I} c_{Ii}\boldsymbol{r}_i - \boldsymbol{R}_I\big).$$

We integrate this by parts, making use of

$$\frac{\partial}{\partial \boldsymbol{R}_I} \delta\big(\sum_{i\in\mathcal{I}_I} c_{Ii}\boldsymbol{r}_i - \boldsymbol{R}_I\big) = -\frac{1}{c_{Ik}} \frac{\partial}{\partial \boldsymbol{r}_k} \delta\big(\sum_{i\in\mathcal{I}_I} c_{Ii}\boldsymbol{r}_i - \boldsymbol{R}_I\big),$$

where we chose an arbitrary atom k involved in the definition of CG site I. If this atom is also involved in the definition of other CG sites, we cannot simply move the partial derivative to $u(\boldsymbol{r}^n)$. In this case, we use a linear combination of all specific atoms j:

$$\frac{\partial}{\partial \boldsymbol{R}_I} \delta\big(\sum_{i\in\mathcal{I}_I} c_{Ii}\boldsymbol{r}_i - \boldsymbol{R}_I\big) = -\sum_{j\in\mathcal{S}_I} \frac{d_{Ij}}{c_{Ij}} \frac{\partial}{\partial \boldsymbol{r}_j} \delta\big(\sum_{i\in\mathcal{I}_I} c_{Ii}\boldsymbol{r}_i - \boldsymbol{R}_I\big),$$

where we require $\sum_{j\in\mathcal{S}_I} d_{Ij} = 1$. Thus we need at least one specific atom per CG site. Integration is then straightforward and yields

$$\boldsymbol{F}_I(\boldsymbol{R}^N) = \left\langle \sum_{j\in\mathcal{S}_I} \frac{d_{Ij}}{c_{Ij}} \left(-\frac{\partial u(\boldsymbol{r}^n)}{\partial \boldsymbol{r}_j} \right) \right\rangle_{\boldsymbol{R}^N}, \tag{3.5.1}$$

where

$$\big\langle \bullet(\boldsymbol{r}^n) \big\rangle_{\boldsymbol{R}^N} \equiv \frac{\int \mathrm{d}\boldsymbol{r}^n \, \exp\big[-u(\boldsymbol{r}^n)/k_\mathrm{B}T\big] \, \delta\big(\boldsymbol{M}^N_{\boldsymbol{R}}(\boldsymbol{r}^n) - \boldsymbol{R}^N\big) \, \bullet(\boldsymbol{r}^n)}{\int \mathrm{d}\boldsymbol{r}^n \, \exp\big[-u(\boldsymbol{r}^n)/k_\mathrm{B}T\big] \, \delta\big(\boldsymbol{M}^N_{\boldsymbol{R}}(\boldsymbol{r}^n) - \boldsymbol{R}^N\big)}$$

is the conditional expectation value (or equilibrium average) in the ensemble of the atomistic system; this expectation value is a function of the CG coordinates $\boldsymbol{R}^N$ (which are kept constant while averaging). In practice, this ensemble average will be replaced by a time average over atomistic trajectories.

Equation (3.5.1) can be broken down into simpler parts. The force on an atom j in

the atomistic system is given by

$$\boldsymbol{f}_j(\boldsymbol{r}^n) = -\frac{\partial u(\boldsymbol{r}^n)}{\partial \boldsymbol{r}_j}.$$

The linear combination

$$\boldsymbol{\mathcal{F}}_I(\boldsymbol{r}^n) = \sum_{j \in \mathcal{S}_I} \frac{d_{Ij}}{c_{Ij}} \boldsymbol{f}_j(\boldsymbol{r}^n) \tag{3.5.2}$$

describes the "atomistic forces acting on CG site I only"; it is the total (collective) force on a CG site I within the atomistic model. (Note that only atoms *specific* to this site contribute.) When CG sites are defined as the centre of mass of a group of atoms, we have $c_{Ii} = d_{Ii} = m_i / M_I$, where $M_I = \sum_i m_i$, and Eq. (3.5.2) simplifies to

$$\boldsymbol{\mathcal{F}}_I(\boldsymbol{r}^n) = \sum_{i \in \mathcal{S}_I} \boldsymbol{f}_i(\boldsymbol{r}^n).$$

The forces $\boldsymbol{\mathcal{F}}_I$ defined by Eq. (3.5.2) are equivalent to the instantaneous collective forces (ICF) that are used, for example, in the Adaptive Biasing Force method. Here we only consider linear combinations of Cartesian coordinates; however, it is possible to derive ICF for arbitrary collective variables, though this leads to significantly more complex equations [109].

The ICF can be considered noisy samples of the mean forces; the optimal CG forces that correspond to a consistent CG model as defined by Eq. (3.3.7) are derived from the atomistic forces through an equilibrium average over the internal degrees of freedom, with the CG coordinates kept constant:

$$\boldsymbol{F}_I(\boldsymbol{R}^N) = \left\langle \boldsymbol{\mathcal{F}}_I(\boldsymbol{r}^n) \right\rangle_{\boldsymbol{R}^N}. \tag{3.5.3}$$

This describes the mean forces on the CG sites. If the CG force field has the flexibility to reproduce Eq. (3.5.3), which will be the case for an N-body potential, this will reproduce the potential of mean force, and the CG model will be consistent with the atomistic model.

3.5.2. Practical Approach: MSCG/FM

Equation (3.5.3) defines the atomistically consistent CG forces. In practice, they need to be approximated by the CG force field $\boldsymbol{G} = \{\boldsymbol{G}_I(\boldsymbol{R}^N)\}$, where $\boldsymbol{G}_I(\boldsymbol{R}^N)$ calculates the force on CG site I as a function of the CG configuration $\boldsymbol{R}^N$. Voth *et al.* introduced a practical approach to the optimisation of $\boldsymbol{G}$, termed the "Multiscale Coarse-

Graining/Force Matching" (MSCG/FM) method [90, 108]. This is based on minimising the residual functional,

$$\chi^2[\boldsymbol{G}] = \frac{1}{3N}\left\langle \sum_{I=1}^{N} \|\boldsymbol{G}_I\big(\boldsymbol{M}_{\boldsymbol{R}}^N(\boldsymbol{r}^n)\big) - \boldsymbol{\mathcal{F}}_I(\boldsymbol{r}^n)\|^2 \right\rangle, \tag{3.5.4}$$

where the average is over the equilibrium canonical ensemble for the atomistic model. Given a sample of K configurations $\Gamma_k = [\boldsymbol{r}^n]_k$ from an all-atom simulation, we can numerically evaluate the residual as

$$\chi^2(\lambda) = \frac{1}{3NK}\sum_{k=1}^{K}\sum_{I=1}^{N} \|\boldsymbol{G}_I(\Gamma_k;\lambda) - \boldsymbol{\mathcal{F}}_I(\Gamma_k)\|^2, \tag{3.5.5}$$

where λ is a set of parameters specifying the CG force field. The MSCG/FM approach assumes that the CG force field is linear in these parameters. (This includes, for example, harmonic approximations to bonded interactions, as well as more flexible potentials such as step functions on a grid or spline interactions, and can even be extended to three-body interactions [26, 27].) Equation (3.5.5) can then be written as a matrix equation, which can be solved in a least-squares sense using QR decomposition. To obtain a smooth potential (splines with continuous derivatives at the grid points), this would require a constrained solver. In practice, the set of all-atom configurations is partitioned into blocks, and an unconstrained solver is used for each block. The results are averaged over all blocks, which leads to an approximately smooth potential.

3.5.3. Connection to Relative Entropy

We can measure the information content in a CG configuration for discriminating between the all-atom model and a CG model with potential U as

$$\Phi(\boldsymbol{R}^N|U) = \log\left[\frac{p_R(\boldsymbol{R}^N)}{P_R(\boldsymbol{R}^N|U)}\right].$$

This can also be considered a configuration-dependent entropy (if both models sample the canonical ensemble):

$$\Phi(\boldsymbol{R}^N|U) = -\beta\big(A(\boldsymbol{R}^N) - U(\boldsymbol{R}^N)\big) + (\text{const.})[U],$$

where the constant may depend on the CG potential. The relative entropy can then be written as

$$S_{\mathrm{rel}}[U] = k_{\mathrm{B}} \int \mathrm{d}^N \boldsymbol{R}\, p_R(\boldsymbol{R}^N) \Phi(\boldsymbol{R}^N | U).$$

Hence, minimising the relative entropy corresponds to minimising the expectation value of Φ. Minimising the residual functional (3.5.4) of the MSCG/FM approach can be shown to correspond to minimising the expectation value of $|\nabla\Phi|^2$ [84]. For a force field that can represent an N-body potential, both approaches will arrive at the potential of mean force. However, when using simpler force fields, for example, using molecular mechanics-style terms such as pair potentials, the two approaches will lead to different effective potentials.

3.6. Summary

In this chapter we compared different approaches to coarse-graining. For "bottom-up" models, we can define a mapping between atomic and CG coordinates. This allows us to define the "consistency" of a CG model in terms of matching the probability distributions of the underlying atomistic model. In momentum space, imposing this consistency condition defines the masses of the CG sites. In configuration space, we find that the consistent CG interaction potential should be equal to the N-body Potential of Mean Force (PMF) for the CG coordinates in the atomistic system.

In practice, current CG models are usually based on pair potentials, and rely on different approaches to find the "best" *effective* pair potential. These approaches can be classified according to whether they are based on matching *structures* or *forces*. Structure-based approaches that aim to reproduce RDFs will all arrive at equivalent pair potentials. In contrast, force-based approaches that aim to approximate the PMF do not necessarily reproduce structural distribution functions. An advantage of force-based approaches is that they rely on *gradients* of the PMF which are defined locally, and hence they can be combined with biasing algorithms such as ABF [66,87]. When a CG approach that aims to match the mean forces in the all-atom model *does* recover structural distributions of the atomistic model, this is strong evidence for a consistent CG model that captures the PMF and will thus be able to reproduce both structural *and* thermodynamic properties, overcoming the representability challenge of coarse-graining. In the next chapter we will introduce a new force-based approach to many-body CG interactions with which we hope to achieve this aim.

4. Gaussian Approximation Potential

All the different approaches discussed in the previous chapter are significantly limited by their common assumption that the CG interactions can be separated into terms similar to the molecular mechanics terms introduced in Section 2.2, such as individual bond and angle terms and non-bonded terms that are described by pairwise additive isotropic site–site interactions ("pair potentials"). This is already an approximation in all-atom simulations [24], where ongoing research is studying extensions such as polarisable force fields [110, 111] and more general many-body approaches [30, 112].* In CG models, only considering molecular mechanics-style terms is a severe approximation to the true interactions defined by the PMF, as each CG interaction is the average of the various interactions in the underlying all-atom model. This is likely to introduce significant correlations between different coarse-grained degrees of freedom, resulting in many-body effects that cannot be captured by simple terms.†

Simple site–site potentials may be able to reproduce radial distribution functions, but this does not imply that they can reproduce higher-order correlations such as angular or orientational distribution functions as well, and such more complex correlations have usually been neglected in previous research.

To obtain a more accurate CG model and, thus, a good description of correlations between different CG sites, we need to include many-body interactions. While there has been some progress in anisotropic potentials [116], and a few forays into three-site interactions [26, 27], demonstrating that three-body interactions are necessary to reproduce the angular distribution function in a one-site water model, previous research on many-body potentials in CG models has been limited. IBI is based only on distance distributions and cannot include any form of more complex correlations. IMC includes

*Many-body potentials such as non-bonded three-body potentials [113], embedded atom models [114], and bond-order potentials [115] have been used extensively for metallic systems for decades. However, in biomolecular modelling, many-body approaches to non-bonded atomic interactions are only now starting to appear.

†With an increasing number of atoms per CG site, the correlations between CG sites are likely to increase as well, and hence the approximation of separating the true interactions into a sum of simple terms, as is typical in coarse-graining, tends to become worse. This is why, historically, (bottom-up) CG models generally grouped just a few atoms within each site.

cross-correlations between different distance distributions; in principle it could take into account three-body interactions, but this would be infeasible in practice [99]. The MSCG/FM approach is based on least-squares fitting and requires interactions to be linear in the parameters; it can be used for general three-body interactions between CG sites, but this has only been applied to one-site descriptions of homogeneous liquids [27], and interactions involving more than three CG sites would likely be infeasible within this approach. The relative entropy approach can in principle handle any functional form for the CG interactions [104], including three-body or more complex interactions, but this has not been applied in practice.

Instead of assuming, as previous approaches, that the PMF can be separated into terms similar to those used in molecular mechanics, the present work is based on a more general many-body approach, the Gaussian Approximation Potential (GAP). GAP is a machine-learning potential based on Gaussian process regression, a type of nonparametric regression with the flexibility to represent, in principle, *any* function. A Gaussian process "learns" a function based on observations of its values and derivatives, allowing us to predict function values at new points even in high-dimensional spaces. Gaussian process regression also comes with the advantage of allowing us to evaluate the *uncertainty* in a prediction. In Section 4.1 we will review Gaussian process regression as a general approach to regression.

GAP was originally developed by Bartók *et al.* [28, 29] in the context of atomistic simulations to infer the potential energy surface from energies and forces obtained from quantum-mechanical calculations such as Density Functional Theory. They demonstrated that GAP can provide nearly the same accuracy as quantum-mechanical methods, in a fraction of the time.

In this thesis we transfer GAP to coarse-grained molecular simulations, with the aim of recovering the many-body PMF $A(\boldsymbol{R}^N)$ for the coordinates of the CG model [Eq. (3.3.7)]. While possible in principle, a direct interpolation of the PMF in $3N$-dimensional configuration space would be impractical. Like the potential energy surface in the atomistic system, the PMF is not just an arbitrary function but has specific properties due to its physical nature. When applying Gaussian process regression to atomic and molecular interactions, it is prudent to take into account our knowledge about the properties of interaction potentials; we will discuss this in Section 4.2. In the context of coarse-graining, the GAP approach faces specific challenges that are less relevant when interpolating a potential energy surface. In Section 4.3 we discuss and develop solutions for the following challenges: noise within force-based coarse-graining, the length scale inhomogeneity of interactions, and sampling issues.

4.1. Gaussian Process Regression

Regression analysis is the study of the relationship $y = f(\boldsymbol{x}) + \varepsilon$ between an input vector $\boldsymbol{x}$ and the observed, potentially noisy output y. Here we make the common assumption that the noise, described by the stochastic variable ε, is independent of the underlying process in the absence of noise, described by the function $f(\boldsymbol{x})$.

We want to infer the unknown function f, given a set of N observations $\{y_n\}$ at positions $\{\boldsymbol{x}_n\}$ (the "training points" or "data"):

$$\mathcal{D} \equiv \{X, \boldsymbol{y}\} = \{(\boldsymbol{x}_n, y_n) \,|\, n = 1, \ldots, N\}.$$

Our aim is to predict the value of the function at a new input position (the "test point") $\boldsymbol{x}^*$. Ideally, we also want to determine the "trustworthiness" of the prediction; we want a measure of the *error* of our predictions. To this end we use Bayesian modelling, a powerful, probabilistic approach to regression analysis.

4.1.1. Bayesian Modelling

In Bayesian modelling, all variables are described by probability distributions rather than a single value. This allows us to explicitly take into account the uncertainty that may generally be present due to approximations, limited sampling, or for any other reasons.

We first consider a parametric model with a fixed functional form $f(\boldsymbol{x}|\boldsymbol{w})$ that depends on a set of parameters $\boldsymbol{w}$. For example, linear regression would be described by the model $f(\boldsymbol{x}|\boldsymbol{w}) = \sum_h w_h x_h$.

We explicitly state our prior assumptions about the functional relationship $f(\boldsymbol{x}|\boldsymbol{w})$, before making any observations, in the distribution $P(\boldsymbol{w})$ over the possible values of the parameters. This is the *prior* distribution. Note that we always have some prior assumptions; making them explicit in the form of probability distributions allows us to reason about them and make more powerful predictions.

The next step is to calculate the *likelihood* $P(\mathcal{D}|\boldsymbol{w})$ of our observations under the model, as a function of the values of the parameters $\boldsymbol{w}$. Finally, we use Bayes's rule to determine the *posterior* probability $P(\boldsymbol{w}|\mathcal{D})$ of the parameters given our observations:

$$P(\boldsymbol{w}|\mathcal{D}) = \frac{P(\boldsymbol{w})\, P(\mathcal{D}|\boldsymbol{w})}{P(\mathcal{D})}.$$

The marginal likelihood $P(\mathcal{D})$ is a constant normalisation factor and does not have to be evaluated directly. Hence, we can write

$$P(\boldsymbol{w}|\mathcal{D}) \propto P(\boldsymbol{w})\, P(\mathcal{D}|\boldsymbol{w}),$$

or schematically:

$$\text{Posterior} \propto \text{Prior} \times \text{Likelihood}.$$

We can predict new function values $f^* = f(\boldsymbol{x}^*)$ by integrating over the parameters:

$$P(f^*|\boldsymbol{x}^*, \mathcal{D}) = \int P(f^*|\boldsymbol{x}^*, \boldsymbol{w}) P(\boldsymbol{w}|\mathcal{D})\ \mathrm{d}\boldsymbol{w}.$$

This is a probability distribution for the unknown function value at the test point $\boldsymbol{x}^*$; we could consider this distribution to be our prediction of the function. In practice, we usually compute the expectation value $\mathbb{E}[f^*]$ to obtain a single value for the prediction. We can compute the variance $\mathbb{V}[f^*]$ as a measure of the "uncertainty" of the predicted value of the inferred function at $\boldsymbol{x}^*$.

The quality of the prediction depends on the choice of model* and the prior distribution of the parameters. Even a probabilistic parametric regression remains constrained by the fixed functional form of the chosen model.

4.1.2. Distributions over Functions

Instead of imposing a functional form and specifying a prior distribution over its parameters, we can directly consider distributions over *functions*. This results in a very flexible, nonparametric model. We can write the posterior probability as

$$P(f(x)|\boldsymbol{y}, X) = \frac{P\big(\boldsymbol{y}|f(x), X\big) P\big(f(x)\big)}{P(\boldsymbol{y}|X)} \propto P\big(\boldsymbol{y}|f(x), X\big) P\big(f(x)\big).$$

We can picture a distribution over functions as an extension to multivariate distributions, in the limit of infinitely many coordinates: consider sampling a function $f(\boldsymbol{x})$ at positions $\boldsymbol{x}_n$, where $n = 1, \ldots, N$. We denote the function values $\{f_n\} = \{f(\boldsymbol{x}_n)\}$ by the vector $\boldsymbol{f}$. Taking $N \to \infty$, we can *define* the function by all the values it takes, $f(\boldsymbol{x}) \equiv \boldsymbol{f}$. Even though functions correspond to vectors of infinite length, in practice there is only ever a finite number of observations and predictions of the function values.

*We could also include the *model* (that is, the specific functional form) in the "parameters". We would then have to specify a prior distribution over models, calculate the posterior, and integrate over all models.

A suitable prior distribution over functions is given by the Gaussian process, a generalisation of the multivariate Gaussian distribution to the infinitely-dimensional vector space of functions. Of course there are many other distributions over functions, but a Gaussian process can be treated analytically rather than having to approximate it numerically, enabling us to handle the "infinite" flexibly.

Whereas a multivariate Gaussian distribution is defined by its mean vector and covariance matrix, the Gaussian process is defined by a mean function $m(\boldsymbol{x})$ and a covariance function $k(\boldsymbol{x}, \boldsymbol{x}')$ (also referred to as the kernel). The defining characteristic of a Gaussian process is that, for any finite subset of input vectors $\{\boldsymbol{x}_n\}$, the function values f_n follow a multivariate Gaussian distribution, $f_n \sim \mathcal{N}(\boldsymbol{m}, \boldsymbol{\Sigma})$, where the mean vector is given by $m_n = m(\boldsymbol{x}_n)$, and the covariance matrix is given by $\Sigma_{nn'} = k(\boldsymbol{x}_n, \boldsymbol{x}_{n'})$. This implies

$$\begin{aligned} m(x) &= \mathbb{E}\big[f(x)\big] \qquad \text{and} \\ k(x, x') &= \mathbb{V}\big[f(x)\big] = \mathbb{E}\big[\big(f(x) - m(x)\big)\big(f(x') - m(x')\big)\big] \end{aligned}$$

for any x, x'.

We can derive the Gaussian process as the extension to infinitely many basis functions of a parametric model that is linear in the parameters. This will be summarised in the following.To keep the notation simple, we show how to model a function of a single scalar argument. However, the derivation straightforwardly generalises to functions over $\mathbb{R}^N$.

We can expand the unknown function f in an arbitrary basis set of H basis functions $\phi_h(x)$:

$$f(x) = \sum_{h=1}^{H} w_h \phi_h(x). \tag{4.1.1}$$

A specific choice of weights $\boldsymbol{w} = \{w_h\}$ then describes a single function. The set of basis functions can be chosen such that we can in principle describe any function.

If we consider the weights w_h to be probabilistic variables, Eq. (4.1.1) describes a prior distribution over functions, determined by our choice of the distribution of w_h. For a set of input points at positions $\{x_n\}$, the corresponding function values $\boldsymbol{f} = \{f_n\}$ are given by

$$f_n = f(x_n) = \sum_h w_h \phi_h(x_n) = \sum_h R_{nh} w_h, \tag{4.1.2}$$

where $R_{nh} = \phi_h(x_n)$ is the $N \times H$ matrix of basis functions evaluated at the input points.

We now make the assumption that the weights w_h are independently and identically distributed according to a multivariate Gaussian distribution with zero mean and constant variance σ_w^2,

$$\boldsymbol{w} \sim \mathcal{N}(0, \sigma_w^2 \mathbb{1}).$$

Then the function values f_n, as linear combinations of Gaussian-distributed random variables, also follow a multivariate Gaussian distribution, again with zero mean. Their covariance matrix Q is given by

$$Q = \langle \boldsymbol{f}\boldsymbol{f}^\mathsf{T} \rangle = \langle R w w^\mathsf{T} R^\mathsf{T} \rangle = R \langle w w^\mathsf{T} \rangle R^\mathsf{T} = \sigma_w^2 R R^\mathsf{T},$$

and hence the function values are distributed according to

$$P(\boldsymbol{f}) = \mathcal{N}(0, Q) = \mathcal{N}(0, \sigma_w^2 R R^\mathsf{T}).$$

Let us now consider a set of radial basis functions of width ℓ, centred at positions h:

$$\phi_h(x) = \exp\left(-\frac{(x-h)^2}{2\ell^2}\right).$$

The matrix multiplication RR^T corresponds to a sum over all basis functions:

$$Q_{nn'} = \sigma_w^2 \sum_h \phi_h(x_n)\phi_h(x_{n'}). \tag{4.1.3}$$

In the limit of infinitely many basis functions, $H \to \infty$, the sum over h becomes an integral. Extending the integration limits to $\pm\infty$ and scaling σ_w^2 appropriately, we can evaluate Eq. (4.1.3) analytically. This results in

$$\begin{aligned} Q_{nn'} &\propto \int_{-\infty}^{\infty} \exp\left(-\frac{(x_n-h)^2}{2\ell^2}\right) \exp\left(-\frac{(x_{n'}-h)^2}{2\ell^2}\right) \mathrm{d}h \\ &= \sqrt{\pi\ell^2} \exp\left(-\frac{(x_n - x_{n'})^2}{4\ell^2}\right). \end{aligned}$$

Hence we can write the elements of the covariance matrix as

$$Q_{nn'} = \delta^2 \exp\left(-\frac{(x_n - x_{n'})^2}{2\theta^2}\right),$$

where we introduced the new length scale $\theta = \sqrt{2}\ell$ and δ^2 contains the prefactor describing the magnitude scale of the covariances between function values.

This means that the covariance matrix is completely described by a covariance *function* between input points,

$$C(x_n, x_{n'}) = \delta^2 \exp\left(-\frac{(x_n - x_{n'})^2}{2\theta^2}\right). \tag{4.1.4}$$

The covariance function is also called the "kernel". Note that there is no longer any dependence on explicit basis functions; the prior is fully specified by the kernel. The covariance function defined by Eq. (4.1.4) is commonly referred to as the squared exponential kernel. This is by no means the only choice, and we discuss the effects of this choice in the next section.

4.1.3. Covariance Function and Hyperparameters

In the previous section we saw that the Gaussian process prior is defined by the covariance function. This relates the covariance between two function values (how much a change in one affects the other) to the "distance" between the corresponding input points: $k(\boldsymbol{x}, \boldsymbol{x}')$ provides a measure of similarity between two input feature vectors $\boldsymbol{x}$ and $\boldsymbol{x}'$. In fact, a distance metric can be defined based on the kernel [117]. This motivates more complex approaches that directly calculate a similarity measure between atomic environments [118]. In principle, any function that results in a positive-definite covariance matrix can be used as a covariance function.

The choice of covariance function determines the "typical behaviour" of functions drawn from the Gaussian process. In this way we can impose properties such as continuity, smoothness/differentiability, periodicity, and so on. Samples drawn from Gaussian processes with different kernels are shown in Fig. 4.1 on the following page.

We now focus on the squared exponential kernel introduced in Eq. (4.1.4). Generalised to d-dimensional input vectors, the squared exponential kernel is given by

$$k(\boldsymbol{x}, \boldsymbol{x}') = \delta^2 \exp\left(-\tfrac{1}{2}\sum_{i=1}^{d}\left(\frac{x_i - x_i'}{\theta_i}\right)^2\right), \tag{4.1.5}$$

where we also introduced the possibility of different length scales θ_i for each input dimension. Equation (4.1.5) is infinitely differentiable, which means that it corresponds to a prior distribution of continuous, smooth functions. This is well suited to potential energy surfaces as they are relatively smooth functions of the atomic coordinates; coarse-graining further smoothens the interaction potential surface. Hence, this is the kernel we use in the remainder of this thesis.

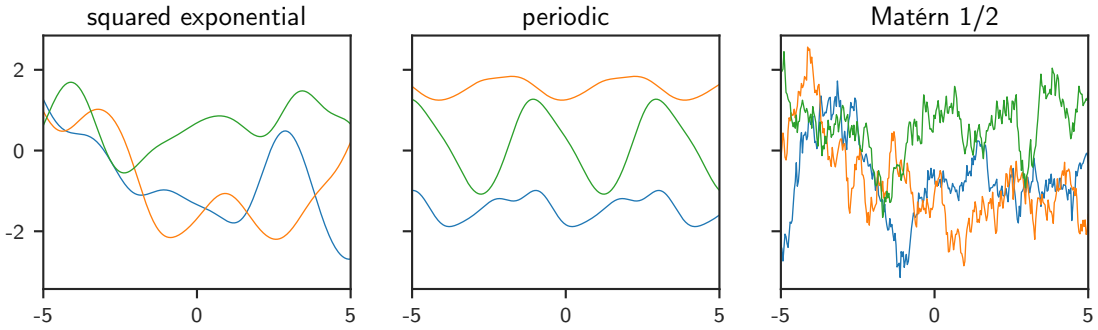

Figure 4.1: Samples from Gaussian processes with different covariance functions. This shows three randomly drawn samples, illustrating different types of priors that can be used in Gassian Process regression. From left to right: squared exponential kernel (4.1.4) with $\delta = 1$ and $\theta = 1$; periodic kernel with period 4: $k(x, x') = \exp[-\frac{1}{2}\sin(\pi(x - x')/4)^2]$; Matérn kernel with $\nu = 1/2$: $k(x, x') = \exp(-|x - x'|)$.

Hyperparameters

Kernels typically depend on *hyperparameters*. This nomenclature distinguishes them from the parameters of a fixed functional form used in a parametric model; the hyperparameters are parameters of the *prior*. The hyperparameters influence the distribution over functions described by the kernel. Our choice of kernel function and hyperparameter values contains our beliefs about the behaviour of the function f; they depend on the system we want to model, but not on the training data.

The squared exponential kernel (4.1.5) has the hyperparameters δ and $\{\theta_i\}$. δ^2 describes the signal variance; hence, δ corresponds to the magnitude of typical fluctuations of the function. θ is the length scale over which the correlation between values of the function at different points decays or, equivalently, the length scale over which the function changes appreciably. This is illustrated in Fig. 4.2(a) on page 54. There can be different values of θ for the different dimensions of the input space.

Ideally, we would marginalise over the space of hyperparameters, for example using Monte Carlo. However, for practical applications this is generally infeasible, and a single value is used for each hyperparameter. In principle it is possible to infer the optimal values of the hyperparameters from the data [32]. However, in the applications we consider, this would be too computationally expensive.

Fortunately, for the molecular systems in which we are interested we can estimate the relevant scales from physical/chemical intuition and other considerations, and hence can determine reasonable values for all the hyperparameters without resorting to automatic determination. This is a great advantage of Gaussian process regression.

Note that the influence of the prior on predictions reduces as the volume of training data increases. With a sufficient volume of training data, the Gaussian process regression is stable with respect to changes in the hyperparameters.

4.1.4. Predictions

As with any distribution, we can draw samples from a Gaussian process; in this case, each sample is a *function*. In practice, we only sample the set of function values on a finite grid. Pictorially, we could consider taking a number of samples from the Gaussian process prior. The mean of these sampled functions approximates the prior mean function and their covariance approximates the covariance function.

We then take into account the observed data: for each sample we evaluate whether it agrees with the data; this corresponds to calculating the likelihood. If a sampled function does not describe the data, we discard it. The set of remaining samples describes the posterior distribution. We can determine the expected function value and its variance at any point by calculating the mean and variance of our samples of the posterior distribution. This is illustrated in Fig. 4.2.

In practice, calculating an accurate mean in this way would require far too many samples to be computationally feasible. Fortunately, for a Gaussian process prior we can determine the mean and variance of the posterior distribution analytically. The posterior distribution will again be a Gaussian process. In the following we show how to predict the function value at a single test point $\boldsymbol{x}^*$.

We consider a Gaussian process prior with covariance function $k(x, x')$ and zero mean function, $m(x) = 0$. This is without loss of generality: we can introduce a fixed mean function $m(x)$ by considering the Gaussian process regression on the modified relation $\tilde{y}(x) = \big(y(x) - m(x)\big)$, which then has zero mean. We can obtain predictions of the original function using $y(x^*) = \tilde{y}(x^*) + m(x^*)$. It is also possible to consider a prior distribution over mean functions and include this in the regression; we refer the interested reader to the literature [32].

The likelihood for the observations $\boldsymbol{y}$ at the N training points $\{\boldsymbol{x}^n\} \equiv X$ is given by the multivariate Gaussian distribution

$$\boldsymbol{y}|X \sim \mathcal{N}(0, C_N),$$

where $C_N = K(X, X)$ is the $N \times N$ matrix of covariances between the training points, with $[K(X, X)]_{mn} = k(\boldsymbol{x}^m, \boldsymbol{x}^n)$.

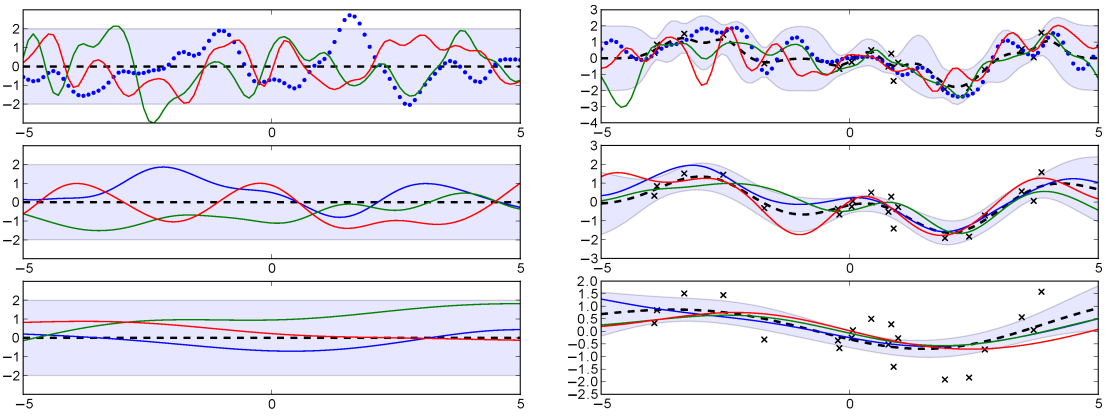

(a) Prior **(b)** Posterior

Figure 4.2: Illustration of Gaussian process regression in one dimension.
(a) Prior distribution using the squared-exponential kernel (4.1.5) [*cf.* Section 4.1.3 on page 51]. The effect of different values for the length-scale hyperparameter is shown, from top to bottom: $\theta = 0.3, 1, 3$. The signal variance is $\delta^2 = 1$. The dashed line denotes the mean function, and the shaded area corresponds to the 95 % confidence interval based on the predicted variance at each point. In each case three samples were drawn at random from the Gaussian process: The dots explicitly show the discrete sampling of a function; the other functions were approximated by joining the sampled points.
(b) Posterior distribution of the Gaussian process after training on the (noisy) data (represented by crosses). The training data was generated by drawing a function from a Gaussian process as in (a) for $\theta = 1$, with additive Gaussian noise of variance $\sigma^2_{\text{noise}} = 0.25$. Note that outside the range covered by the data, the variance of the posterior distribution returns towards that of the prior.

We now consider the joint distribution under the Gaussian process prior of observing both the training data $(X, \boldsymbol{y})$ and the function value f^* at a test point $\boldsymbol{x}^*$. This is again a multivariate Gaussian distribution:

$$\begin{bmatrix} \boldsymbol{y} \\ f^* \end{bmatrix} \Bigg| X, \boldsymbol{x}^* \sim \mathcal{N}\left(0, \begin{bmatrix} K(X,X) & \boldsymbol{k}(X,\boldsymbol{x}^*) \\ [\boldsymbol{k}(X,\boldsymbol{x}^*)]^\intercal & k(\boldsymbol{x}^*,\boldsymbol{x}^*) \end{bmatrix}\right), \tag{4.1.6}$$

where $\boldsymbol{k}(X, \boldsymbol{x}^*)$ is the vector of the covariances between the test point $\boldsymbol{x}^*$ and each of the training points. The joint probability for this multivariate Gaussian distribution is given by

$$P(f^*, \boldsymbol{y}) \propto \exp\left(-\tfrac{1}{2}\begin{bmatrix} \boldsymbol{y}^\intercal & f_* \end{bmatrix} C_{N+1}^{-1} \begin{bmatrix} \boldsymbol{y} \\ f^* \end{bmatrix}\right), \tag{4.1.7}$$

where C_{N+1} is the covariance matrix of the joint probability,

$$C_{N+1} = \begin{bmatrix} [C_N] & [\boldsymbol{k}] \\ [\boldsymbol{k}^\intercal] & [\kappa] \end{bmatrix},$$

with $\kappa = k(\boldsymbol{x}^*, \boldsymbol{x}^*)$. We could calculate C_{N+1}^{-1} by brute-force matrix inversion, but in fact the inverse can be calculated more elegantly [119]. Writing the inverse as

$$C_{N+1}^{-1} = \begin{bmatrix} [M] & [\boldsymbol{m}] \\ [\boldsymbol{m}^\intercal] & [m] \end{bmatrix}$$

and setting $C_{N+1}C_{N+1}^{-1} = \mathbb{1}$, we obtain

$$\begin{aligned} m &= (\kappa - \boldsymbol{k}^\intercal C_N^{-1}\boldsymbol{k})^{-1}, \\ \boldsymbol{m} &= -mC_N^{-1}\boldsymbol{k}, \\ M &= C_N^{-1} + \frac{1}{m}\boldsymbol{m}\boldsymbol{m}^\intercal. \end{aligned}$$

We now want to determine the *conditional* distribution for f^*, given $\boldsymbol{x}^*$ and the training data $(\boldsymbol{y}, X)$. This amounts to marginalising over all other rows in Eq. (4.1.7).

For multivariate Gaussian distributions this can be calculated exactly [31, 32]. The result is our posterior distribution:

$$f^*|\boldsymbol{y}, X, \boldsymbol{x}^* \sim \mathcal{N}(\hat{f}_*, \sigma_*^2),$$

corresponding to

$$P(f^*|\boldsymbol{y}) \propto \exp\left(-\frac{(f^* - \hat{f}_*)^2}{2\sigma_*^2}\right),$$

where the expectation value or mean of the prediction is

$$\hat{f}_* = \boldsymbol{k}^\intercal C_N^{-1} \boldsymbol{y} \tag{4.1.8}$$

and the variance, describing the uncertainty in the mean, is

$$\sigma_*^2 = \kappa - \boldsymbol{k}^\intercal C_N^{-1} \boldsymbol{k}. \tag{4.1.9}$$

$\hat{f}_*$ as a function of $\boldsymbol{x}^*$ corresponds to the mean function of the posterior Gaussian process. While the mean function of the prior was assumed to be zero, the mean function of the *posterior* distribution is not zero. From now on we will identify the result of our prediction with the expectation value, $f^* \equiv \hat{f}_*$.

An important aspect of Gaussian process regression is that we only need to calculate C_N^{-1}, and not the inverse of C_{N+1}. This allows us to separate the calculation into a "training" stage involving the matrix inversion (scaling cubically with the number of training points N), and a prediction that scales linearly in N. We can write the expectation value (4.1.8) of the predicted posterior distribution as

$$f^* = f(\boldsymbol{x}^*) = \sum_{n=1}^{N} \alpha_n \, k(\boldsymbol{x}^n, \boldsymbol{x}^*), \tag{4.1.10}$$

where the sum is over all training points. The coefficients α_n are given by

$$\{\alpha_n\} \equiv \boldsymbol{\alpha} = C_N^{-1} \boldsymbol{y}.$$

This corresponds to a linear combination of the kernel function centred on the training data points, similar to kernel ridge regression [32]. Hence, in contrast to parametric regression, the functional form can be considered to change with every new observation. This is what allows the nonparametric regression using Gaussian processes to, in principle, interpolate *any* function.

Noisy observations

So far we assumed exact, noise-free observations. However, in practice, observations are generally afflicted with some error. As mentioned at the beginning of Section 4.1,

we typically consider the noise to be independent and additive to the "true" functional relationship $f(\boldsymbol{x})$. We further assume that the error in observations follows a Gaussian distribution with variance σ^2_{noise}. This can be described by a conditional probability of observations $\boldsymbol{y}$ on the true function values $\boldsymbol{f}$:

$$P(\boldsymbol{y}|\boldsymbol{f}) \sim \mathcal{N}(\boldsymbol{f}, \sigma^2_{\text{noise}}\mathbb{1}),$$

where the mean is given by the exact function values, and the standard deviation σ_{noise} describes the noise level.

We can account for this in the Gaussian process regression by including a noise term in the covariance between observations at training points. Instead of being given directly by the kernel function, the covariance between two observations in the training set is now given by

$$[C_N]_{mn} = \text{cov}(y^m, y^n) = k(\boldsymbol{x}^m, \boldsymbol{x}^n) + \sigma^2_{\text{noise}}\delta_{mn}.$$

Note that the kernel function only depends on the *positions* of the input points, whereas the noise term only depends on the *indices*. This captures the assumption that two independent observations at the same position will still differ due to noise. As we do not want to add noise to the predictions, the covariance between training points and test point remains unchanged. Equivalently, we can add the noise covariance to the noise-free kernel matrix:

$$C_N = K + \sigma^2_{\text{noise}}\mathbb{1}.$$

The prediction of the Gaussian process regression for the function value at a test point is then given by

$$\hat{f}_* = \boldsymbol{k}^\intercal[K + \sigma^2_{\text{noise}}\mathbb{1}]^{-1}\boldsymbol{y}. \tag{4.1.11}$$

Here we assumed the same noise covariance for all training points. However, we can also consider a diagonal noise covariance matrix with different noise levels for each training point. Each observation can have its own noise level.

This way of accounting for noisy observations is the Bayesian approach to regularisation. In contrast to other approaches such as the MSCG/FM method (even when it is based on Bayesian statistics [120]), Gaussian process regression provides a *physical* noise parameter that has the units of the physical observables. This means that we can directly specify the uncertainty of the training data in terms of physical quantities; we can set the actual noise level which we may be able to determine from the training data.

Notes on variance prediction

It is a key feature of Bayesian modelling that we consider *distributions* over variables, and that we do not obtain just a single prediction value but can also determine, for example, the variance of the predicted distribution, indicating how strongly we believe our prediction.

In Gaussian process regression, the predicted variance is given by Eq. (4.1.9). Note that this variance only depends on the *positions* of training and test points, not on the values of the predicted function. In practice, the predicted variance at a given test point is just a measure of the "distance" from the training points. Thus, the value of the predicted variance may need to be treated with caution; for example, if the noise parameter is set too small, the Gaussian process regression will predict a very small variance near the training points, even though this is just a result of overfitting. Hence it is important to remember that the predicted variance only gives an indication of whether the test point is in a region that is sufficiently well-covered by the training data.

Nevertheless, the predicted variance can be very useful for diagnosing failures of Gaussian process regression and to find out when our model has encountered a part of the input space that was not covered by the training set.

In practice, the calculation of the predicted variance involves a small regularisation constant to ensure that the matrix inversion of is well-behaved. Because of this, in the limit of high data coverage, the predicted variance tends towards this regularisation constant. This will become apparent when considering the GAP-predicted variance in Chapter 5/Fig. 5.4. In the limit of zero data coverage, the predicted variance tends towards the variance of the prior, δ^2.

4.1.5. Sums and Derivatives

In our application of Gaussian process regression to coarse-graining, we need to be able to train from and predict derivatives and sums of function values.

Derivative predictions

We can straightforwardly predict derivatives of the inferred function value by taking the derivative of Eq. (4.1.8) with respect to coordinates x_i^* of the test point:

$$\frac{\partial \hat{f}^*}{\partial x_i^*} = \frac{\partial \boldsymbol{k}^\intercal}{\partial x_i^*} C_N^{-1} \boldsymbol{y},$$

where the partial derivative on the right-hand side is the derivative of the covariance function. For the squared exponential kernel (4.1.5), this is given by

$$\frac{\partial k^n}{\partial x_i^*} = \frac{x_i^n - x_i^*}{\theta_i^2} k(\boldsymbol{x}^n, \boldsymbol{x}^*).$$

Derivative observations

We can also train the Gaussian process from derivative observations [121, 122]. Differentiation is a linear operator and hence the derivative of a Gaussian process is again a Gaussian process.

The covariance between a function value observation and a derivative observation is given by

$$\operatorname{cov}\left(y^m, \frac{\partial y^n}{\partial x_i^n}\right) = \frac{\partial C(\boldsymbol{x}^m, \boldsymbol{x}^n)}{\partial x_i^n},$$

and the covariance between two derivative observations is given by

$$\operatorname{cov}\left(\frac{\partial y^m}{\partial x_i^m}, \frac{\partial y^n}{\partial x_j^n}\right) = \frac{\partial^2 C(\boldsymbol{x}^m, \boldsymbol{x}^n)}{\partial x_i^m \partial x_j^n}.$$

We can simply extend the covariance matrix C_N accordingly, and include the derivative observations in the vector $\boldsymbol{y}$ for Eq. (4.1.8). Being able to consistently include derivative observations is a significant advantage of Gaussian process regression.

Sum observations

Within the Gaussian process regression we can also learn from observations of sums of function values [28]. We can describe the linear combination of function values as

$$F_m = \sum_n L_{mn} f(x_n).$$

Returning to Eq. (4.1.2) on page 49, we now write

$$F_m = \sum_{n,h} L_{mn} R_{nh} w_h.$$

This still follows a Gaussian distribution, but now there is a new covariance matrix given by

$$Q' = \langle \boldsymbol{F}\boldsymbol{F}^\intercal \rangle = \langle LRww^\intercal R^\intercal L^\intercal \rangle = \sigma_w^2 LRR^\intercal L^\intercal = LQL^\intercal,$$

where $Q = \langle \boldsymbol{f}\boldsymbol{f}^{\mathsf{T}} \rangle$ is the covariance matrix for the "hidden" function we want to infer from observations of sums of its values. A similar derivation can be carried out for observations of sums of derivatives. Predictions of sums are trivially obtained by adding the predictions of the individual terms.

4.1.6. Sparsification

The linear algebra operations (specifically, the matrix inversion) involved in teaching the Gaussian process and determining the α vector scale as N^3 in computing time and as N^2 in memory, where N is the number of training observations. Hence the number of training points that can be included in a regression is limited by the available computational resources. (The prediction at a test point only scales in time as N for the mean and as N^2 for the variance.)

In Section 4.2 we will see that in our application of Gaussian process regression the training data has a large amount of redundancy: there will be many similar input points, but as we only learn from *total* forces on all CG sites, we need to include all training points and cannot simply thin out the training set. This would give rise to values of N beyond computational tractability.

When the input points are highly correlated, we can use *sparsification* to reduce the computational cost without sacrificing much accuracy. Instead of including all closely spaced points, we introduce S latent pseudo-inputs $\{\bar{\boldsymbol{x}}_s\}$ ($s = 1, \ldots, S$), the "sparse points", that support the function we want to learn.

Instead of the function predictions being directly related to the observations and covariances between training points, we have a "hidden layer" of sparse points, and we can consider this procedure as two sequential Gaussian process regressions: in the first one we obtain the prediction at the sparse points based on the original training points, and in the second one we obtain the prediction at the test point based on the latent values at the sparse points.

There are various different approaches to the sparsification approximation [123]. The implementation in GAP relies on the Deterministic Training Conditional (DTC) approximation.* We calculate the covariances between sparse points:

$$[K_{SS}]_{ss'} = k(\bar{\boldsymbol{x}}_s, \bar{\boldsymbol{x}}_{s'}),$$

*The more recently developed Fully Independent Training Conditional (FITC) approximation by Snelson and Ghahramani [124] is theoretically a better approximation, and was originally used in GAP [28]. However, in practice it was found that FITC is much slower than DTC when including derivative observations in the training set, whereas their predictive performance is comparable.

and between sparse and input points:

$$[K_{NS}]_{ns} = k(\boldsymbol{x}_n, \bar{\boldsymbol{x}}_s).$$

The latent observations $\bar{\boldsymbol{y}}$ at the sparse points are given by

$$\bar{\boldsymbol{y}} = K_{NS}{}^{\intercal}\Lambda^{-1}\boldsymbol{y},$$

where $\Lambda = \sigma^2_{\text{noise}}\mathbb{1}$ is the noise variance. The covariance matrix between the latent observations at the sparse points is then given by

$$Q_S = K_S + K_{NS}{}^{\intercal}\Lambda^{-1}K_{NS}.$$

This is the sparse covariance matrix used in the predictions. In predicting the function value at a test point $\boldsymbol{x}^*$, we now require the covariance between the test point and the sparse points:

$$[\bar{\boldsymbol{k}}]_s = k(\bar{\boldsymbol{x}}_s, \boldsymbol{x}^*).$$

Finally, the sparse predictions for function value and predicted variance are given by

$$\begin{aligned}\hat{f}_* &= \bar{\boldsymbol{k}}^{\intercal}Q_S^{-1}\bar{\boldsymbol{y}} \qquad \text{and} \\ \sigma^2_* &= k(\boldsymbol{x}^*, \boldsymbol{x}^*) - \bar{\boldsymbol{k}}^{\intercal}(K_S^{-1} - Q_S^{-1})\bar{\boldsymbol{k}}.\end{aligned}$$

This reduces the required computing time to NS^2 for the teaching and to S and S^2 for prediction of mean and variance, respectively. As $S \ll N$, this significantly reduces the computational cost.

In our application, the sparse points are chosen as a subset of the full set of training points. We will discuss different approaches to the selection of sparse points in Section 4.2.5.

4.2. Interaction Potentials

In Chapter 3 we saw that the "true" interactions for a consistent CG model are given by the potential of mean force (PMF) $A(\boldsymbol{R}^N)$ for the CG coordinates $\boldsymbol{R}^N$. Our aim is to learn the coarse-grained interaction potential from derivatives of the PMF.* This

*In principle, it is possible to measure the value of the PMF directly, for example, using Metadynamics. However, this would require a long time to convergence, is not as accurate as methods based on integrating derivative observations, and current applications of approaches such as metadynamics are generally

can be based directly on observations of mean forces determined by methods such as Constrained or Restrained MD [65]. Alternatively, we can infer the PMF by averaging over the instantaneous collective forces [66]. In principle, we could use Gaussian process regression to infer the PMF directly as a function of the coordinates of all the CG sites, $A(\boldsymbol{R}_1, \ldots, \boldsymbol{R}_N)$, from the observations of mean or instantaneous collective forces. However, this would not be a practicable approach for interaction potentials, as in coarse-graining the PMF is very high-dimensional. Moreover, we would be neglecting the knowledge we have about the properties of physical systems, such as symmetries: for instance, the PMF is invariant under translations and rotations of the whole system and under permutations of equivalent sites or molecules. In this section we discuss how to take into account this knowledge. This distinguishes the Gaussian Approximation Potential (GAP) from a direct application of Gaussian process regression.

4.2.1. Locality Approximation

The exact N-body PMF depends, in principle, on the positions of all the CG sites in the system. In practice, molecular systems are *localised*: The mean force on any particular site only depends on the relative positions of nearby sites, and does not vary significantly under perturbations of other sites that are sufficiently far away.

This can be demonstrated explicitly in a "locality test". For a given site or molecule, we determine all others within a certain cutoff radius; this constitutes the "neighbourhood" of the central site/molecule. We keep the positions of all CG sites within the neighbourhood fixed, and measure the variance of the mean force(s) on the central site/molecule as we vary the positions of the CG sites outside the neighbourhood. As the neighbourhood is made larger, this variance decreases. We shall demonstrate this in Section 5.3.2.

The locality of the mean forces motivates a separation of the N-body PMF $A(\boldsymbol{R}^N)$ into a sum of local contributions,

$$A = \sum_{\xi} \sum_{i \in \xi} w^{\xi}(\mathcal{N}_i), \tag{4.2.1}$$

where ξ describes different types of local contributions, i iterates over all instances of this type in a given configuration of the system, and $\mathcal{N}_i$ describes the "neighbourhood" or local environment on which the i^{th} local contribution depends. Such a separation

limited to no more than six collective variables.

into local terms is a feature of *all* interatomic potentials.* Instead of inferring the full PMF, Gaussian process regression can then be used to infer the local free energy contributions w^{ξ}.

The definition of the neighbourhood $\mathcal{N}_i$ depends on the type ξ, but it can be considered to contain all or a subset of the positions of the CG sites within a certain cutoff radius; for example, this can be one central site and all other sites within a cutoff distance (illustrated in Fig. 4.3 on the following page), or the sites describing a pair of molecules whose centres of mass are within a cutoff.

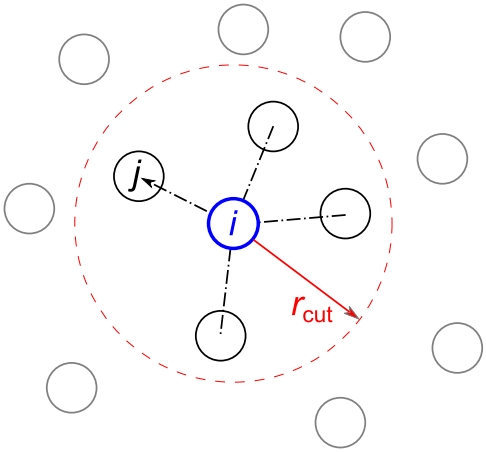

Figure 4.3: The neighbourhood $\mathcal{N}_i$ of a site i (blue, centre). Only sites within a cutoff radius R_{cut} (dashed red circle) are included. The neighbourhood consists of the set of (vectorial) distances $\boldsymbol{R}_{ij} = \boldsymbol{R}_j - \boldsymbol{R}_i$ to the other sites j (dash-dotted lines) for which $|\boldsymbol{R}_{ij}| < R_{\text{cut}}$.

To make use of this approach in Gaussian process regression, we still need to specify how the local free energy contribution w^{ξ} depends on the neighbourhood. We note that the molecular mechanics-style interaction terms commonly used in coarse-graining approaches can all be described within this framework, though this is often not discussed explicitly. There, ξ would specify the different types of bonds, angles, dihedral angles, and pairs of non-bonded sites. However, in those cases the "neighbourhood" is very limited.

For example, the neighbourhood for a CG bond between two sites i and j would consist of just these two sites. The CG bond would then be associated with a contribution to the PMF given by $w^{\text{bond}}(\boldsymbol{R}_i, \boldsymbol{R}_j)$. But the physical properties of the bond do not depend on its orientation in space; the dependence of the contribution to the free

*There may be additional long-ranged terms, such as electrostatic and/or dispersion forces, that do not fall into this schema and would need to be treated explicitly.

energy needs to be invariant under the symmetries of the system. We could impose this symmetry on the local free energy contribution by considering $w^{\text{bond}}\big(|\boldsymbol{R}_i - \boldsymbol{R}_j|\big)$. This motivates the definition of a function $d(\boldsymbol{R}_i, \boldsymbol{R}_j) = |\boldsymbol{R}_i - \boldsymbol{R}_i|$ that converts the positions of the sites within a neighbourhood into an argument to a local interaction potential term, while incorporating invariance under translation, rotation, and permutation of sites. We call this a "descriptor", in the following denoted $\boldsymbol{d}^\xi$. More generally, this allows us to write Eq. (4.2.1) as

$$A = \sum_\xi \sum_{i \in \xi} w^\xi \Big(\boldsymbol{d}^\xi \big(\{ \boldsymbol{R} \in \mathcal{N}_i \} \big) \Big).$$

All interaction terms can be considered as a combination of a local free energy contribution and a descriptor. For the molecular mechanics-style interaction terms, the descriptor value is one-dimensional, and the local free energy term could be fitted just as well by a spline. However, in general the descriptors may return a multidimensional vector.

The choice of descriptors used in modelling a system corresponds to our assumptions about how the PMF can be separated into various local contributions. As discussed throughout this thesis, only considering a separation into one-dimensional descriptors such as non-bonded site–site interactions can be too restrictive for accurate CG simulations.

For molecular systems, a much less restrictive approximation is a separation of the PMF into terms corresponding to whole molecules, to pairs of molecules, and to triplets of molecules. This amounts to assigning an internal free energy contribution to each monomer configuration, and free energy contributions to all two- and three-body interactions, where "body" refers to the *entirety* of a molecule (not just one of its individual sites).

Here we have neglected long-range electrostatic and dispersion/van-der-Waals interactions. These could be treated separately, either by including them in the fitting procedure or by setting partial charges based on the underlying atomistic system. But in biological and other solvated systems, these long-range interactions are often screened by the solvent, and in the following we will not consider explicit long-range interactions. In this work, we will only include their contribution to the effective short-range interactions within the locality cutoff.

The forces (instantaneous collective forces or mean forces) that we can observe are gradients of the *total* potential of mean force A. The Gaussian process can be trained on sums of derivatives and thereby infer the "hidden" local free energy contributions w^ξ

(which do not have a direct physical representation). Together with the compactness of the neighbourhoods, this is what allows us to apply Gaussian process regression to larger systems, for which learning the PMF $A(\boldsymbol{R}^N)$ directly would be infeasible. As the local free energy contributions only depend on local neighbourhoods, the number of interactions only grows linearly with the number of CG sites in the system. This allows GAP to be a viable approach to coarse-graining. In the following section we discuss the descriptors that we use in the present work in more detail.

4.2.2. Descriptors

A descriptor is a function that transforms the neighbourhood defined in terms of the Cartesian coordinates of CG sites into a (smaller) set of variables, the "features", taking into account relevant symmetries. These features are then used as the input positions in the Gaussian process regression. In our approach, each descriptor is associated with a local free energy contribution, which is a function of the descriptor value.

Descriptors can be one-dimensional scalars, such as the distance between two sites, or a set of values, such as the three distances defining a three-site monomer, or more generally the set of all distances between CG sites that describe a monomer or a pair (dimer) or triplet (trimer) of molecules. Descriptors can also represent more complex features such as an expansion of the relative positions of all sites within the neighbourhood into spherical harmonics.

A descriptor is the representation of certain degrees of freedom of a system. Hence, the descriptor values should be invariant under transformations that do not change physical aspects of the system. Specifically, the local interaction terms have to be invariant under rotations and translations of the set of all coordinates that are used in a descriptor.

Descriptors should be "unique": similar neighbourhoods should correspond to similar descriptor values, and distinct neighbourhoods should correspond to distinct descriptor values. Together with the covariance function, descriptors have to provide a similarity measure. Physically equivalent neighbourhoods, which differ at most by a transformation that leaves the local free energy contribution unchanged, should map to the same descriptor values.

We now discuss the descriptors that will be used later in this thesis. This includes the site–site distance descriptor that is equivalent to splines or tabulated pair potentials used in many previous approaches to coarse-graining as well as *molecular* descriptors that describe entire molecules (monomers) and their combinations as pairs and triplets.

Pairwise site–site distance

GAP based on the pairwise distance descriptor is equivalent to the commonly used pairwise splines between sites. This descriptor only depends on two sites at a time and simply returns the distance between the two sites. It is useful for describing the repulsive core potential between sites and as an independent verification that GAP can generate potentials that are equivalent to those obtained by the MSCG/FM approach.

Monomer-based terms

In our work, we focus on *molecular* rather than site-based descriptors, with local contributions to the free energy associated with each single molecule (internal or monomer free energy), pairs of molecules (two-body or dimer interaction), and triplets of molecules (three-body or trimer interaction). It is important to distinguish these molecular terms from site-based terms; for example, for a three-site CG description, the molecular dimer term already corresponds to 6-body interactions between CG sites.

The molecular descriptors can be considered as a cluster expansion of the PMF:

$$A(\boldsymbol{R}_1,\ldots,\boldsymbol{R}_N) = A(\boldsymbol{R}_{\{i,\alpha\}}) \approx \sum_j w_{\text{mono}}(\boldsymbol{R}_{j,\{\alpha\}}) + \sum_{j,k} w_{\text{dimer}}(\boldsymbol{R}_{j,\{\alpha\}},\boldsymbol{R}_{k,\{\beta\}}) + \sum_{j,k,l} w_{\text{trimer}}(\boldsymbol{R}_{j,\{\alpha\}},\boldsymbol{R}_{k,\{\beta\}},\boldsymbol{R}_{l,\{\gamma\}}),$$

where Roman letters index molecules and Greek letters index sites within each molecule. Here we only consider systems with a single type of molecule; however, the descriptors are not limited to such systems, and we could just as easily consider, for example, a trimer descriptor for three different types of molecules.

The monomer descriptor considers all monomers in the system. The dimer and trimer descriptors consider pairs and triplets of monomers that are within a cutoff based on the centre of mass distances. In each case, the descriptor is constructed from the list of all pairwise distances between the sites involved in a monomer, pair, or triplet:

$$\{d_{ij} \mid i = 1,\ldots,n;\ j = i+1,\ldots,n\}$$

where $d_{ij} = |\boldsymbol{R}_j - \boldsymbol{R}_i|$. The descriptor dimension D is given by the number of distances, $D = k(k-1)/2$, where k is the number of sites involved in a descriptor. For a homogeneous system, the descriptor dimension as a function of the number of CG sites per molecule is given in Table 4.1.

CG description of molecule	descriptor dimension		
	monomer	dimer	trimer
one site	—	1	3
two sites	1	6	15
three sites	3	15	36

Table 4.1: Dimensions of molecule-based descriptors for one-, two- and three-site CG descriptions.

Note that this can be an over-complete description, with the number of distances included in the descriptor being larger than the number of degrees of freedom (which is given by $3k - 6$, where we exclude the "external" degrees of freedom corresponding to overall translation or rotation of the monomer/dimer/trimer). For example, a two-site trimer is described by 15 distances but only has 12 degrees of freedom.* This does not affect the regression.

It would be appealing to base the molecular descriptors directly on certain degrees of freedom, such as the centre-of-mass (COM) distance and orientation of the monomers with respect to each other, instead of including all intermolecular distances. However, it is impossible to describe orientations in such a way that the descriptor has continuous derivatives everywhere. Because of this, the interaction potential surface in the descriptor space would not be smooth and hence this could not be captured well by the Gaussian process regression.

In a one-site model, there is no monomer term, and the dimer term is equivalent to a pairwise spline. Together with a simple three-site term, this is already described well in the MSCG/FM approach. GAP becomes particularly relevant for descriptions of molecules with several CG sites, where higher-order terms are not well represented by combinations of site-based interactions. However, the dimensionality of the feature space grows very quickly with the number of sites. Hence, GAP will be primarily useful when we want to describe individual residues with a comparatively small number of sites, as in Ref. 88, where one lipid molecule of 144 atoms is described by 3 to 5 sites.

A tetramer descriptor would be computationally infeasible for several reasons. First, the number of descriptor dimensions increases factorially with the number of sites; we would require 28 distances for a two-site model or 66 for a three-site model. Not only does this increase the time taken for each evaluation of the covariance, but we

*It is also possible to consider only a subset of all the distances, though we do not make use of this feature in the present work.

would also need many more training and sparse points to accurately describe the interaction potential in this high-dimensional descriptor space, making both the GAP training and predictions much more costly. Second, in larger systems there can be many more combinations of four molecules than of pairs or triplets, vastly increasing the number of descriptors that are found in each configuration, and slowing down the force calculation even further.

Other approaches

Based on the observation that the kernel in the Gaussian process regression measures the "similarity" between two input vectors, there are other approaches such as SOAP [118] that, instead of combining a descriptor with the squared exponential kernel, directly calculate the similarity between two environments based on all positions within a cutoff radius. This has been shown to work very well for recovering local atomic energies.

However, preliminary investigations quickly showed that the use of molecule-based terms made it much easier to reproduce the mean forces. This suggests that it is advantageous to associate local free energy contributions with whole molecules and pairs/triplets of molecules rather than with individual CG sites, whereas SOAP would associate a local free energy with each CG site. We note that SOAP might be useful in a solvent-free CG description of ions, where hydration effects and the polarisation of the water close to the ions lead to three-body effects [125].

Covariance cutoff

To ensure that the number of descriptor instances only grows linearly with the number of sites in the model, each descriptor has an explicit cutoff and depends at most on the positions of sites within that cutoff. The cutoff distance may be measured from a central site, or between centres of mass; it is not necessarily a site–site distance.

We explicitly impose locality on the descriptors by multiplying the covariance by a cutoff function. We use a cutoff function based on the centre-of-mass distance between monomers; this cutoff function goes to zero as the distance approaches the cutoff radius. Hence the covariance for descriptors beyond the cutoff will be exactly zero. To ensure smooth behaviour of the derivatives (forces), the cutoff function smoothly goes from one to zero over a certain transition width; in practice, a transition width of at least 3 Å results in smooth forces.

The cutoff is directly related to the locality of the potential of mean force. In practice, typical values for the descriptor cutoffs are 3 Å to 12 Å. For monomer-based descriptors,

the cutoff can be chosen based on minima of the COM RDF. For dimer interactions, this may extend to the second or third minimum. For trimer interactions, the first minimum is usually sufficient. Increasing the cutoff will generally increase the resulting accuracy, at a higher computational cost. Hence, we can check whether a cutoff is large enough by comparing results with those using an increased cutoff. If there is no significant difference, the smaller cutoff was sufficient.

Permutational symmetry

The features given by the descriptors, such as the ordered list of all pairwise distances, already incorporate the invariance under rotation or translation of the group of molecules involved. However, this may not yet account for the invariance of a local free energy contribution under permutations of sites. For example, a dimer of two identical monomers with two different types of sites is invariant under exchange of the monomers. If the sites within a monomer are of the same type, then the descriptor is also invariant under exchange of the two sites within one monomer.

For small descriptor dimensions, permutational symmetry can be imposed directly within the descriptor. For larger numbers of descriptor values, it is easier to enforce the invariance under permutations at the level of covariances. Given the unsymmetrised covariance $k(\boldsymbol{d}^m, \boldsymbol{d}^n)$, we can turn this into a symmetric covariance by summing over the set $\mathcal{P}$ of all permutations that leave the system invariant:

$$k^p(\boldsymbol{d}^m, \boldsymbol{d}^n) = \sum_{p \in \mathcal{P}} k\big(\{d^m_{1,2,\ldots,n}\}, \{d^n_{p_1,p_2,\ldots,p_n}\}\big).$$

4.2.3. Core Potential

When building a new interaction potential using GAP, it is possible to use any pre-existing potential as a "core potential". This corresponds to giving the Gaussian process prior a fixed mean function. In practice, the forces predicted by the core potential for the training points are subtracted from the observed forces, and a Gaussian process is trained on these differences. In the final prediction the results from core potential and difference-taught GAP are added together again. There are two distinct applications for core potentials.

First, this allows us to make use of previous results, for example, if we already have a good monomer potential and now want to create a dimer potential. Introducing the monomer potential as a core potential for the training of the dimer descriptor reduces the variance or "ruggedness" of the energy landscape that has yet to be described.

Hence, the function that needs to be inferred is smoother, which makes it easier to learn and can increase the accuracy of the overall potential.

Second, a core potential can be used to cover regions for which there is otherwise not sufficient training data to accurately represent the free energy surface in that region. This will likely be the case for configurations with high free energy, typically containing CG sites at short range due to the strong short-range repulsion between the atoms involved. The issue of insufficient training data will be discussed further in Section 4.3.3.

4.2.4. Multiple Descriptors

In practical applications, it may not be possible to provide separate training sets for different descriptors such as monomers and dimers. This may be due to technical reasons, or simply because it may not be meaningful to isolate monomers or dimers. For example, a lipid molecule can be expected to have very different properties and interactions in water and inside a lipid membrane. For this it is important that within the GAP approach we are able to teach multiple descriptors simultaneously from the same training data.

Different descriptors are independent of each other. Hence, the covariance between local free energy contributions associated with different descriptors is always zero. This means that the full covariance matrix will be block-diagonal, with one block per descriptor type. The relative values for the δ variance hyperparameters that describe the scale of fluctuations of functions determine how much of the variance of the target function ends up being assigned to the different descriptors.

For example, when describing a system of three molecules, the trimer descriptor has the same flexibility for describing free energy surfaces by itself as when it is combined with monomer and dimer descriptors. Hence, we need to set the ratios of the δ hyperparameters according to the fraction of the free energy contribution that we expect to be due to monomer, two-body and three-body terms. The advantage of including monomer and dimer contributions explicitly is that, in our applications, the largest contribution to the forces on CG sites is due to monomer terms; the monomer term alone already predicts most of the force on a site. This means that the left-over contribution to the forces, due solely to higher-body interactions, is smaller: this corresponds to a successive flattening out of the part of the free energy surface that is left to predict, which makes it easier to fit the trimer term, thus requiring less data and fewer sparse points.

4.2.5. Selection of Sparse Points

When we discussed sparsification of Gaussian process regression in Section 4.1.6, we assumed the set of sparse points as given. We now discuss methods for constructing this set. Given the potentially high dimensionality of the space of descriptor values and the large number of training and sparse points we use in GAP, the selection of sparse points should, ideally, be automatic, without requiring user intervention.

The chosen sparse points need to be able to accurately represent the input data as given by the training points. Predictions become less accurate the further away a test point is from the nearest sparse points. Hence, the sparse points need to cover that subspace of descriptor values that is *accessible* by the model. Note that there are large parts of the space of descriptor values that will have no training points at all and will never be relevant to a simulation. This can be for two reasons; firstly, because the descriptor may be overcomplete, hence some values of the descriptor cannot correspond to any configuration (for example, because a set of distances does not fulfil the triangle inequality), and secondly, because descriptor values can correspond to configurations that are completely unphysical, such as two overlapping molecules. To accurately represent the functional relationship of the local free energy contributions, there should be a higher density of sparse points where the function value changes faster. We now present a range of different sparsification methods. Their performance in practical applications will be compared at the end of Chapter 5.

Random selection

The simplest but least effective sparsification method is a random selection of sparse points from the entire training set. This selection is uniform over training points, and hence directly equivalent to the distribution of training data. If we simply take unbiased samples from an all-atom simulation (which will be the case for learning from ICF as in the MSCG/FM method), this corresponds to a Boltzmann-weighted distribution. This results in a large number of sparse points in the minima of the PMF where the mean forces (MF) are actually zero; much fewer sparse points will correspond to configurations of molecules at short range for which the MF are large and repulsive. For this reason, random selection may not provide enough sparse points at the "edges" of the training set.

Uniform selection

Instead of a random selection of sparse points, which is sensitive to the distribution of the training points, we can attempt to create a uniform selection across the whole descriptor space.

For this we construct a d-dimensional grid, where d is the number of dimensions of the descriptor. The number of bins per grid dimension is given by $n = \lceil S^{1/d} \rceil$, where S is the number of sparse points. The bins in each grid dimension are evenly distributed between the minimum and maximum value in that dimension across all training points. The training points are then binned into this grid, and one point selected from each bin that is not empty. If there are less than S non-empty bins, a non-empty bin is selected at random and another point from this bin added to the set of sparse points; this process is repeated until S sparse points have been collected.

This technique is better at making sure there is an even distribution of sparse points; this has the advantage that it avoids selecting points that are close to each other, which would give rise to a high condition number of the covariance matrix. If there can be enough bins per grid dimension, uniform selection works best; this is certainly the case in few dimensions. However, uniform selection may struggle in higher dimensions. For the 15-dimensional 3-site dimer, there are at most two separate bins per dimension for any number of sparse points less than $2^{15} = 32768$. Interestingly, though, uniform selection of sparse points still performed best even in this case.

CUR decomposition

CUR decomposition can be used as a low-rank matrix approximation of an $m \times n$ matrix A, similar to Singular Value Decomposition (SVD). SVD decomposes the matrix A into a product of three matrices: the $m \times m$ unitary matrix of left-singular vectors of A (eigenvectors of $AA^\intercal$), the $m \times n$ diagonal matrix of singular values, and the $n \times n$ unitary matrix of right-singular vectors of A (eigenvectors of $A^\intercal A$).

The CUR decomposition approximates $A \approx CUR$, where the $m \times c$ matrix C contains c columns of A and the $r \times n$ matrix R contains r rows of A. The $c \times r$ matrix U is defined as $U = C^\intercal A R^\intercal$. This is not a unique decomposition; in the present work, we implemented the algorithm by Mahoney and Drineas [126], which aims to select those columns and rows from the data that result in the best low-rank fit of the data matrix.

CUR decomposition is less accurate than SVD, but easier to understand: Whereas the left and right singular vectors represent the data in a rotated space and do not have a straightforward interpretation, the selected rows in R directly correspond to those

data points that lead to the best approximation of the data matrix. In the sparsification procedure, we take advantage of this by determining which training points are most relevant to approximate the full $N \times D$ matrix of descriptor values for the whole set of training points. These points are then used as the sparse points. Our current implementation of CUR decomposition does not take into account the permutational symmetry of descriptors; however, this does not seem to significantly impede the performance of this selection method in practice. CUR decomposition may be more appropriate than uniform selection specifically for higher-dimensional descriptors or when the number of sparse points is comparatively low.

The selected columns in C correspond to those features that best distinguish between data points; while we do not make use of this yet, this could be used to reduce the number of descriptor dimensions.

Covariance-based selection

It is also possible to "greedily" select sparse points directly based on the covariance function. For this, we consider a "working set" of sparse points, initially filled with a single, arbitrarily chosen point from the set of training points. We then calculate the predicted variance for all other training points. (Note that the variance predicted by the Gaussian process regression only depends on the *positions* of observations, not the values.) We add to the working set that point which has the largest variance, that is, the point whose function value would be least explained by the working set of sparse points. This procedure is iterated until we have selected the desired number of sparse points.

While this approach may provide a good selection of sparse points, its disadvantage is that the calculation of the "intermediate" variance takes increasing long as the size of the working set grows. For a large number of sparse points, this selection procedure becomes prohibitively slow.

A related, but significantly more efficient approach to selecting sparse points has recently been published by Schreiter *et al.* [127]. Implementing and testing this approach in the context of GAP remains the subject of future work.

4.3. Application to Coarse-Graining: GAP-CG

As discussed previously, the Potential of Mean Force (PMF) depends on the state point, and this carries over to the local free energy contributions that are associated with the descriptors. For example, the dimer interaction between two molecules will depend on their environment: we cannot expect a vacuum dimer potential to accurately describe the bulk model. We need to train GAP descriptors in the same environment in which we want to use them. However, due to the locality of the mean forces, it should be possible to use the local free energy functions trained on data from small all-atom bulk systems to subsequently simulate the CG model in a larger volume of the same density, using the local free energy contributions as "building blocks" for the PMF of larger systems. The typical workflow for using the Gaussian Approximation Potential in Coarse-Graining (GAP-CG) is illustrated in Fig. 4.4.

When using GAP to obtain interactions for coarse-grained models, observations of the function values are generally not available. GAP can be trained solely on derivatives, and the Gaussian process regression automatically incorporates the integration of the derivative observations to the free energy surface of the CG system. However, this brings up several challenges that are not encountered as prominently when developing interatomic potentials for which we can easily evaluate the total energy. Here we discuss these challenges, relating to noise and averaging in the GAP approach to coarse-graining (Section 4.3.1), how to deal with inhomogeneous length scales (Section 4.3.2), and how to deal with issues resulting from insufficient sampling (Section 4.3.3).

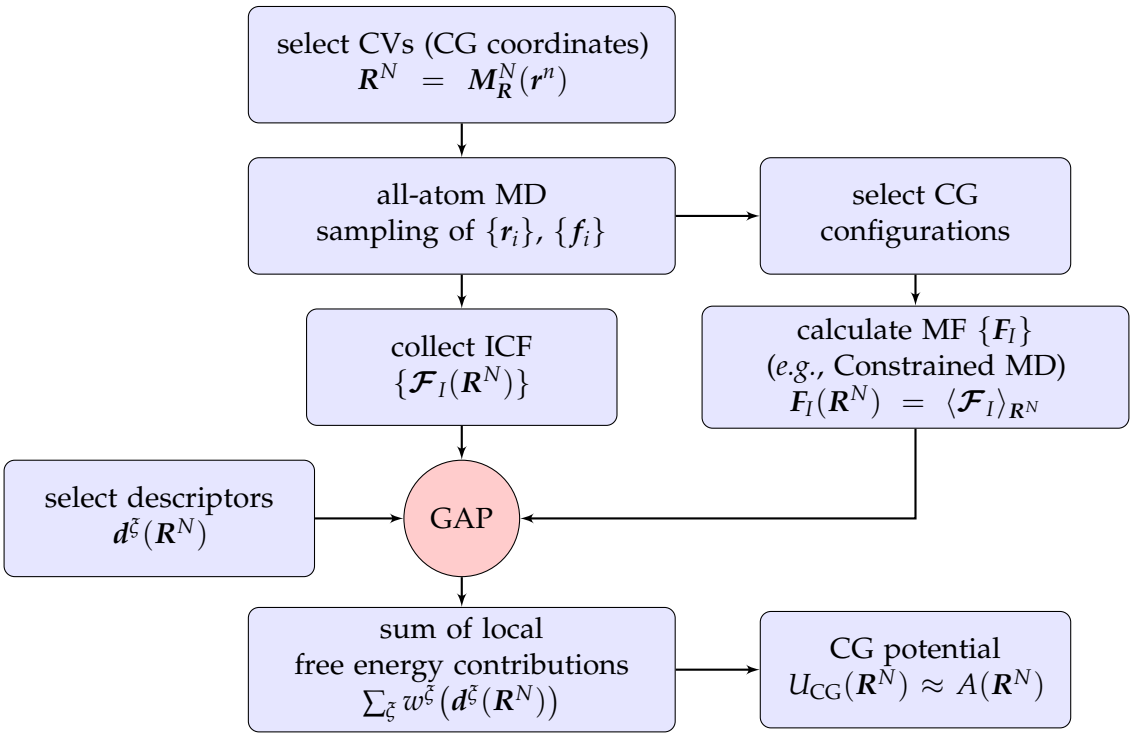

Figure 4.4: Workflow for GAP-CG

4.3.1. Noise

Gaussian process regression is an inherently regularised approach that can automatically embody Occam's Razor [128]: we should always consider the simplest model that describes our observations. Correspondingly, a smoother function that does not overfit the observations is better at generalising to unknown test points.

Using the noise level hyperparameter σ_{noise}, we can tune the Gaussian process regression between closely interpolating the training data and averaging noisy observations. Note that we can directly specify the noise level as a *physical* quantity; this is a significant advantage over other approaches. In many cases, other methods such as least squares minimisation also include some kind of regularisation such as Tikhonov regularisation; however, the available regularisation parameter usually does not have any direct physical interpretation. The Voth group also investigated a Bayesian extension to their MSCG/FM method [120]; however, their approach relies on a precision hyperparameter without direct physical meaning that needs to be optimised numerically, resulting in a high computational cost.

We can teach a CG GAP from mean forces that have been explicitly obtained by methods such as Constrained or Restrained MD. When possible, this is the favoured approach, as it is always easier to learn from data with smaller error. Such methods typically provide the estimated uncertainty in the resulting mean forces. When learning from instantaneous collective forces (ICF), the noise level must be set appropriately so that GAP can correctly infer the underlying average. In principle, it is possible to measure the noise level in the training data by binning the ICF and measuring the variance in each bin. Note that this binning refers to the *descriptor space*; this will not be considered in this thesis except for the simplest one-dimensional case, and determining the noise level from ICF for more complex descriptors remains for future work.

To prevent overfitting, it is generally better to use a larger noise level hyperparameter than one that is too small. When a test set of mean forces is available, the noise level can be gradually reduced from above until the prediction error on the test set no longer decreases. However, this can only be relied on when the test set is completely statistically independent of the training set and at the same time contains a sufficient sample of configurations. Hence, in practice, it is preferable not to "optimise the noise level" and to use a somewhat larger noise level hyperparameter instead.

It is important to note that, apart from the stochastic noise in derivative observations, there are also systematic errors involved in using the descriptors to represent the PMF. As discussed in the previous section, we aim to reproduce the PMF using a sum of

local free energy contributions that are associated with descriptors. With the exception of small model systems, for which it is possible to use descriptors with the same dimensionality (for example, a trimer cluster described by the trimer descriptor), this is an approximation; note that this applies to *all* coarse-graining approaches that aim to reproduce the PMF. This inherent (and unknown) approximation is given by the discrepancy between the actual shape of the free energy surface in the $3N$-dimensional configuration space and the kinds of hyperdimensional surfaces that can be represented by the descriptors. The mean forces are the gradients of the *actual* PMF, but it may be impossible to describe the exact PMF as a function of all CG coordinates using the available descriptors. Hence, with respect to the "approximate representation" of the PMF using descriptors, even fully-converged mean forces with negligible stochastic error can still contain noise when considered from the point of view of the descriptors. This needs to be taken into account when setting the noise level for the Gaussian process regression.

As descriptors are generally unable to *exactly* describe a system, they only capture an approximation of the full PMF. This means that the choice of training data can bias the resulting interaction potential; this is the case, for example, when the training set contains both trimer and tetramer configurations. Hence, the training set needs to consist of unbiasedly sampled configurations consistent with the environment which the CG potential shall describe.

4.3.2. Inhomogeneous Length Scales

In many cases, the length scale parameter θ is assumed to be constant throughout the whole range of a descriptor coordinate. In the present work, we focus on descriptors that are based on distances between sites. However, the interactions between atoms — and, by extension, CG sites — do not have a single length scale. Typically, in the long-range regime at large intermolecular separations, the forces due to two- or higher-body interactions are relatively weak and slowly changing, and eventually go to zero. In contrast, at short range the electrostatic repulsion of overlapping electron clouds leads to a strong and quickly increasing repulsion. This means that the length scale of the interaction changes as a function of the descriptor values.

If we use a constant length scale appropriate to the behaviour at short range, $\theta = \theta_{\mathrm{SR}}$, then at larger distances we will overfit the noise. Moreover, we will lose predictive power, as this reduces the correlation length and hence more sparse points are needed to sufficiently cover the descriptor space.

If we use a constant length scale appropriate to the behaviour at long range, $\theta = \theta_{\mathrm{LR}}$, we induce correlations between force observations that are actually very different from each other; this results in the strong short-range repulsion leaking over into longer range and leads to oscillations ("ringing") in energies and forces. This effect becomes increasingly large as the value of θ is increased. The ringing can be reduced by increasing the noise level hyperparameter, at the cost of reducing the overall prediction accuracy.

The problem of inhomogeneous length scales in Gaussian processes was discussed by Gibbs [129]. One way to address this issue would be to introduce a non-stationary covariance in which the length scale depends on the values of the descriptors. However, this would require significant changes to and an increase of the complexity of the code base. A much simpler and just as effective solution to the problem is to introduce a *transfer function* that stretches a coordinate before presenting it to the Gaussian process regression. We can thus "bend down" the steep short-range repulsion, thereby effectively increasing the scaled distance over which the function varies at short range, increasing the *effective* length scale at short range and thereby bringing it closer to the length scale at long range. This results in a more uniform effective length scale in the transformed coordinate. The transfer function is much more straightforward to implement than a non-stationary covariance, as the core implementation of Gaussian process regression (such as the covariances) does not need to be changed.

To minimise the number of additional hyperparameters, we make the assumption that the length-scale changes of different distances within the descriptor are independent of each other. This allows us to use a one-dimensional transfer function on each coordinate separately. However, we can still use different transfer functions for the coordinates; for example, we would want to distinguish between bonded and non-bonded distances.

In the following, we consider a single descriptor coordinate x. Introducing the transfer function $T(x)$ leads to a new effective descriptor $\tilde{x}$, given by the transformation $x \to \tilde{x} = T(x)$. Without the transfer function, the aim of the Gaussian process regression is to infer the local free energy term $w(x)$. With the transfer function, the Gaussian process has to learn a different function in the transformed input space, $\tilde{w}(\tilde{x})$. The free energy should not depend on the transfer function and, hence, we require

$$w(x) = \tilde{w}\big(T(x)\big).$$

This implies a transformation of the gradients $\mathrm{d}w/\mathrm{d}x$ of the free energy term:

$$\frac{\mathrm{d}w}{\mathrm{d}x} = \frac{\mathrm{d}\tilde{w}}{\mathrm{d}\tilde{x}} T'(x),$$

where $T'(x) = \mathrm{d}T/\mathrm{d}x$ is the derivative of the transfer function and $\mathrm{d}\tilde{w}/\mathrm{d}\tilde{x}$ is the derivative observation used in the Gaussian process regression. To ensure that the inferred free energy contributions remain smooth, the first derivative of the transfer function must be continuous.

We now consider the functional form of $T(x)$. Our aim is to have one, longer length scale at long range and another, shorter (effective) length scale at short range, and a smooth transition from one to the other. Note that a multiplicative factor on the distance will stretch the axis as a whole. Hence, we want an "effective" multiplicative factor that depends on the distance and smoothly changes from 1 in the long-range part to the scaling factor $\alpha = \theta_{\mathrm{LR}}/\theta_{\mathrm{SR}}$ at short range. This multiplicative factor corresponds to the derivative of the transfer function, and an appropriate effective factor is given by the tanh function:

$$\frac{\mathrm{d}T}{\mathrm{d}x} = (1-\alpha)\tfrac{1}{2}\left[\tanh\left((x-r_0)/w\right)+1\right] + \alpha. \tag{4.3.1}$$

For $x \ll r_0$, this goes to α, and for $x \gg r_0$, this goes to 1. We can obtain the actual transfer function by integrating Eq. (4.3.1). This results in

$$T(x) = (1-\alpha)\tfrac{1}{2}\left[w \log\left(2\cosh((x-r_0)/w)\right) + x + r_0\right] + \alpha x. \tag{4.3.2}$$

This transfer function depends on three parameters: r_0 is the mid-point for the transition between the length scales, w is the extent of the transition region either side of r_0, and α is the stretching factor.

In the present work, we apply the transfer function only to the non-bonded distances in the dimer descriptor.

4.3.3. Sampling Issues

All coarse-graining approaches that are based on sampling of an underlying system encounter problems when the sampling is insufficient. This is particularly apparent in methods that rely on learning interactions only from derivative observations.

We can distinguish two types of insufficient sampling that are relevant to coarse-graining. The first one is the failure to sample a completely different state of the

system when there are other free energy basins that, due to free energy barriers, were not encountered during the sampling of the underlying all-atom model. Addressing this problem is the focus of various free energy methods such as ABF sampling [72] and Metadynamics [74, 75], and will not be considered further in this work. The second type of insufficient sampling, discussed in the following, corresponds to insufficient sampling of configurations with large free energy. Specifically, this refers to neighbourhoods or local environments associated with descriptors with large local free energy contributions.

The probability of finding the atomistic system in a certain configuration is given by the Boltzmann factor $\exp(-\beta U)$, where U is the potential energy. Thus, configurations with a high potential energy — often due to the highly repulsive core interaction between two (or more) atoms that are very close to each other — are exponentially suppressed. Such atomic configurations map to CG configurations that should similarly have a low probability, corresponding to a large free energy.

The strong core repulsion between atoms (due to the electrostatic repulsion between overlapping electron clouds) means that configurations in which atoms are at close range are highly penalised; the system spends very little time in such configurations. This carries over to the CG description: for two CG sites to be close to each other, the atoms involved need to be close to each other. Correspondingly, there are few CG configurations that contain descriptors at short range. Hence, there will be few or even zero samples of regions of the descriptor spaces that are associated with large free energy contributions. In these regions, the CG interaction should likewise be strongly repulsive. However, in an absence of training data, the Gaussian process regresses to the mean of the prior, *i.e.*, zero. Because of this regression to zero, as two CG sites are brought closer to each other, the free energy will reach a maximum and then return to zero once the corresponding descriptor ends up outside the sampled region. Hence, instead of providing a strong short-range repulsion, the predicted mean force will decrease in magnitude and eventually turn attractive, trapping the system in an unphysical minimum. One could employ longer length scale hyperparameters in the hope that this would provide sufficient correlation with strongly repulsive force observations at distances where there is still reasonable sampling. However, this would lead to the problem of artificial "ringing" of the force magnitude over the whole range of the descriptor as discussed in the previous section.

When inferring the CG interactions from mean forces, we could improve sampling by choosing the CG configurations from high-temperature simulations, as long as the mean forces are evaluated in simulations at the target temperature. However, this

only *reduces* the likelihood that the CG model encounters undersampled short-range configurations, and is not feasible when collecting instantaneous collective forces.

To systematically overcome the issues arising from insufficient sampling in regions of high free energy, we need to provide a core potential that assigns a large free energy to configurations with CG sites at close range, so that the CG system is repelled and cannot fall into an unphysical minimum beyond the initial short-range repulsion.

The easiest and most appropriate core potential is a repulsive site–site pair potential with a short-range cutoff. As this is described by a one-dimensional descriptor, it is easy to provide sufficient training data, there needs to be only a small number of sparse points, and it is easy to visualise the resulting potential to ensure its validity. As the system only rarely visits configurations with sites at close range, and given the small number of sparse points needed to accurately describe one-dimensional potentials, including such a repulsive short-range core potential adds no significant computational cost.

While we could choose a fixed functional form (such as a Lennard–Jones potential) and appropriately set its parameters, we can also obtain the short-range repulsion from the available training data. For example, we can determine the repulsive short-range potential by fitting a monomer descriptor together with a pairwise distance descriptor with a short-range cutoff. Then the pairwise distance descriptor will pick up only the (repulsive) short-range contribution to the interactions between non-bonded CG sites. These pairwise short-range interactions can be visualised and manually adjusted beyond the coverage by training data to ensure that they provide an appropriate short-range repulsion. We then use this repulsive short-range potential as a core potential for teaching the full interaction potential.

4.4. Summary

In this chapter we gave an overview of Gaussian process regression and how it can be combined with local descriptors to derive interaction potentials; this combination constitutes the Gaussian Approximation Potential (GAP). We transferred this approach to CG interactions inferred from derivative observations of the Potential of Mean Force (PMF). Using an appropriate choice of descriptors, we can infer local contributions to the free energy surface; these are "building blocks" of the PMF for molecular CG models.

GAP can infer the PMF directly from observations of either mean or instantaneous collective forces. In this regard, GAP-CG is similar to the MSCG/FM approach.

However, the descriptors available in GAP provide for a much more flexible, powerful approach. GAP only depends on very few settings or hyperparameters, most of which have a physical meaning and can be related to our prior knowledge about systems. Moreover, GAP allows us to quantify our confidence in the predictions, immediately highlighting issues related to a lack of coverage by training data.

In the following chapter, we will present the results of applying GAP-CG to molecular systems, demonstrating that we can recover the PMF to a high accuracy. We will also discuss the influence of the hyperparameters on the accuracy of the resulting potentials.

5. GAP-CG: Recovering the Potential of Mean Force

In Chapter 3 we showed that a CG model is consistent with the underlying atomistic potential when the CG potential U_{CG} is equal to the many-body PMF $A(\boldsymbol{R}^N)$, a function of the coordinates of all N CG sites. A full N-body potential is generally infeasible, so the practical challenge is to find a way to represent the PMF.

In Chapter 4 we introduced the Gaussian Approximation Potential (GAP) as a general, flexible many-body approach to interaction potentials. GAP infers the interaction potential from observations of the forces acting on the sites in the model. In the GAP-CG approach we make the assumption that the PMF can be separated into a sum of *molecular* terms, assigning local free energy contributions to internal configurations of molecules (monomer terms) and to dimers and trimers of molecules.

In this chapter we show results of applying GAP-CG to various test systems. These results support our assumption that the PMF can be separated into a sum of molecular terms and demonstrate that this leads to much more accurate CG models than the more common approximation of separating CG interactions into terms based on pairwise radial interactions between CG sites. To this end, we compare GAP to approaches relying on pair potentials, both structure-based (*cf.* Section 3.4) and force-based (*cf.* Section 3.5).

We note that studies of coarse-graining approaches tend to analyse only radial distribution functions (RDFs) and, in a few cases, angular distribution functions (ADFs).* This can provide a reasonable characterisation for CG models in which each molecule is described by a single site.†

In the following, we consider the more interesting case of molecules that retain internal structure in their CG description and, hence, have a defined orientation. As

*Ref. 130 is a rare example of a paper that considers distributions that are more complex than RDF and ADF, in this case the distribution of dipole orientation *vs.* distance.

†For one-site CG models, the pairwise splines between sites correspond to molecular two-body interactions, and Larini *et al.* showed that explicitly including three-body interactions in a one-site model of liquid water "can dramatically improve the reproduction of structural properties of the original atomistic system" as measured by RDF and ADF [26].

an example, we show results for a two-site description of methanol and a three-site description of benzene. We discuss their chemical composition and the CG mapping in Section 5.1.

Including the orientation within the CG description allows us to consider more complex distributions to characterise the structure of the system than just RDF and ADF. Structure-based approaches recover the RDF by construction. However, reproducing the RDF does not imply that the model will also reproduce more complex structural properties such as collective variables of the CG system, or correlations between them. In fact, we will see that this is not the case, in general. Force-based approaches have no guarantee of reproducing the structure. While this can be a severe limitation for pair potentials, we will see that the more complex descriptors available in the GAP approach can reproduce the structure better than even structure-based approaches.

Using the trimer descriptor, GAP can *exactly* represent the full N-dimensional PMF of systems containing up to three molecules. In Section 5.2 we demonstrate that GAP is able to recover the PMF from force observations for small methanol clusters (monomers, dimers, and trimers). In contrast, CG models based on pairwise radial interactions fail to describe the overall structure correctly — even if they reproduce individual distance distributions to a high accuracy.

For larger systems, GAP would still be exact *if* the PMF actually separates into a sum of molecular monomer/dimer/trimer terms. In Section 5.2.4 we demonstrate that this is a much better approximation for the methanol tetramer than only considering pairwise site-based interactions.

Clusters are just toy systems to demonstrate the GAP approach; there is no practical need for coarse-graining small systems. Coarse-graining only becomes relevant for large systems. Hence, we also consider the application of GAP-CG to bulk systems. We show that in bulk the decomposition of the PMF into contributions due to pairs and triplets of molecules provides a very good approximation to the results of all-atom calculations. In bulk, a large contribution to the effective interactions is simply due to the packing structure. Even a potential that only models the core repulsion can already qualitatively reproduce various structural distributions. This makes it easier for structure-based approaches such as IBI to reproduce not just the RDF but also other distributions. However, this does not mean that such pair potentials accurately represent the CG interactions. We show that GAP, based on matching forces, reproduces the RDF as well as the pair potentials obtained by the structure-based IBI approach and even provides a *better* description of the correlation between relative orientation and distance at large molecular separations. We discuss this in Section 5.3, where

we consider both methanol and, to show that our approach can easily be applied to different systems, benzene.

In Section 5.4 we discuss how hyperparameters, sparsification and the noisiness of the training data influence the generation of CG GAP potentials.

5.1. Model Systems: All-Atom Description and CG Mapping

We use methanol and benzene as the test systems for comparing GAP and competitive CG approaches. In this section we briefly discuss these molecules and their properties.

5.1.1. Methanol

Methanol is a small molecule with the chemical formula CH_3OH. It is the simplest alcohol, consisting of a methyl group (CH_3) and a hydroxy group (OH). Its structure is shown in Fig. 5.1.

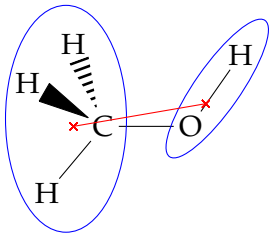

Figure 5.1: Structure of the methanol molecule. The blue outlines indicate the grouping of atoms to CG sites; the CG coordinates are given by the centre of mass of all the atoms in a group and are indicated by red crosses.

The oxygen in the hydroxy group is considerably electronegative and carries a partial charge of $q_O = -0.65e$. This is primarily taken from the hydrogen directly bound to it, which ends up with $q_{H_O} = +0.42e$. Methanol molecules can form hydrogen bonds between their hydroxy groups, significantly influencing the structures they form. Each methanol molecule can be an acceptor and a donor of a hydrogen bond, so each molecule tends to form two hydrogen bonds.

CG mapping

In previous research on the coarse-graining of liquid methanol, each molecule has usually been represented by a single site [94, 108], neglecting all information about the

orientation of the molecules.* Here, we will use a two-site description, with one site at the centre of mass of the methyl group and a second site at the centre of mass of the hydroxy group.

Note that this does not allow us to model the highly directional nature of the hydrogen bond explicitly, as this is determined by the position of the hydrogen atom relative to the oxygen–carbon bond. However, we will see that GAP-CG can recover the hydrogen-bonded structure of the methanol even without an explicit description of the hydrogen bond direction.

5.1.2. Benzene

Benzene (C_6H_6) forms an aromatic six-membered carbon ring, with one hydrogen attached on the outside of each carbon atom. Due to its delocalised π orbitals, this molecule has a very rigid planar structure. The structure of benzene is illustrated in Fig. 5.2.

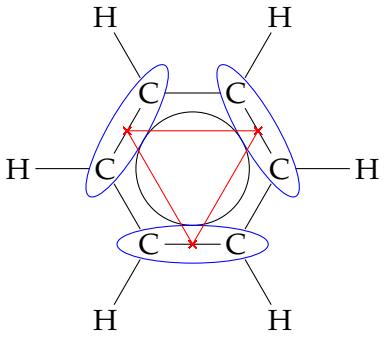

Figure 5.2: Structure of the benzene molecule. The blue outlines indicate the grouping of atoms to CG sites; the CG coordinates are given by the midpoints between pairs of carbon atoms and are indicated by red crosses. The hydrogen atoms are not included in the CG description.

CG mapping

To be able to represent the orientation of the planar benzene molecule, we employ a three-site CG description.† We place one CG site on the centre of mass for each

*In Ref. 107, Izvekov and Voth employ a two-site description equivalent to the one used in this work, but they only compare site–site RDFs with the all-atom model.

†As a further simplification, we could approximate the six-fold rotational symmetry with continuous rotational symmetry and describe a benzene molecule with two sites along the plane normal axis. However, this could not be considered within the linear mapping framework of Eq. (3.2.2) and would require more

pair of neighbouring carbon atoms. This choice is consistent with two different CG descriptions for each molecule in an all-atom configuration, though they are physically equivalent. Note that our CG description only has three-fold rotational symmetry, whereas the benzene ring actually has six-fold symmetry. In Section 5.3.4 we will show that GAP is able to recover the six-fold symmetry of the all-atom interactions in the CG model. The hydrogen atoms are not part of our CG description. Their influence is only captured through the forces they exert on the carbon atoms.

5.1.3. Simulation Protocol

All simulations were carried out at $T = 300\,\mathrm{K}$. At this temperature, methanol and benzene are in the liquid phase, which is relevant to a wide range of applications in biology and chemistry. For the all-atom simulations, we use the AMBER MD package [131], with the ff99SB (Hornak & Simmerling) force field [36]. Letif Mones refined the force field for benzene using GAUSSIAN [132] for geometry optimisation and calculating the partial charges. To ensure accurate MD simulations, we choose a time step of 0.5 fs for all atomistic simulations.

It is important to note that here and in the following the aim is not to reproduce real-world properties of the systems we discuss, but to investigate how well CG models can reproduce the properties of the all-atom model on which they are based.

Unless specified otherwise, we derive the force-based potentials from mean forces, so that we can distinguish issues relating to the averaging procedure from issues arising from the approximation of the PMF as a sum of molecular terms. We will discuss how GAP can include the averaging and be trained on instantaneous collective forces in Section 5.4.3. As discussed in Section 4.3.3, we are specifically interested in obtaining mean forces where their magnitude is large, and hence all mean force calculations in this thesis are based on Constrained MD, for which we use the PMFlib toolkit [63]. For methanol clusters, the Constrained MD simulation for each CG configuration was run for 10 ns. For bulk systems, Constrained MD for all CG coordinates is much more expensive and mean forces were collected in runs of only 100 ps.

Coarse-grained simulations were run with QUIP [133] (the in-house code for GAP [29, 118]*) and LAMMPS [134] (for pair potentials).

complicated calculations to obtain the forces on the CG sites.

*The GAP code is available for non-commercial use from `www.libatoms.org`.

5.2. Methanol Clusters

When the PMF can be separated into a sum of molecular terms described by the descriptors, GAP can recover the full PMF [as defined by Eq. (3.3.7)]. This is proven by the CG MD distributions being indistinguishable from the CG view of the all-atom distributions. This will certainly be the case when a single descriptor can completely characterise the system configuration. With the trimer descriptor this is the case for up to three molecules. Here we consider clusters of one to four methanol molecules.

By considering in turn monomer, dimer, and trimer, we can use the potential from the smaller system as the core potential for the next one. This means we only need to learn genuine two-body effects in the dimer, and only genuine three-body effects in the trimer. This allows us to optimise the parameters for each descriptor separately, simplifies the modelling process and makes better use of available training data.

For a cluster of four methanol molecules, we are no longer able to reproduce the PMF exactly. However, we will see that GAP based on molecular descriptors is a much better approximation to the all-atom simulation than approaches based on pair potentials.

5.2.1. Monomer

In our two-site CG model of methanol, the monomer has a single (internal) degree of freedom, the CG bond length. This allows us to visualise all properties of the model, which is no longer possible for more complex systems. Moreover, this makes it very easy to capture the CG interaction.

In our CG description, the CG bond between the CH_3 and OH sites is very similar to the carbon–oxygen bond in the underlying all-atom simulation. The mean distance is somewhat larger, as the inclusion of the hydrogen atoms shifts the centre of mass away from carbon and oxygen atom positions. The CG force–distance relation is almost linear. Whereas the carbon–oxygen bond in our all-atom model is a perfectly harmonic spring, the CG bond also contains contributions from other interactions. For example, bending of the angles will slightly shift the CG bond length, leading to a (very small) anharmonicity. Nevertheless, the CG bond is described well by a harmonic approximation, and the corresponding force constant is comparable to the C–O all-atom force constant of $23.98\,\mathrm{eV/\AA^2}$.

CG potentials

We can obtain the CG bond potential in various ways, both structure-based and force-based, which we now describe.

Structure-based (Boltzmann Inversion). As there are no other degrees of freedom, Eq. (3.4.1) accurately describes the distribution of the bond length, and Boltzmann Inversion is, in principle, exact. Inaccuracies are due to finite sampling. We derive the Boltzmann-inverted potential from 2.6×10^6 samples of the CG bond length in the all-atom model. The probability distribution is determined by a histogram of the bond lengths with 100 bins between 1.36 Å and 1.65 Å. Equation (3.4.2) then gives the CG bond potential. In order to calculate the derivative, we fit a cubic spline to the resulting potential.

The bond force is almost perfectly linear, corresponding to a Gaussian-distributed CG bond length, and the bond potential is well described by the approximation of a harmonic potential. We calculate mean $r_0 = \mathbb{E}[p(r)/r^2]$ and variance $\sigma_r^2 = \mathbb{V}[p(r)/r^2]$ of the Jacobian-corrected distribution of CG bond lengths. In the harmonic approximation, the CG bond potential is given by $U(r) = \frac{1}{2}k(r - r_0)^2$, with the force constant $k = k_\mathrm{B}T/\sigma_r^2$. For our model, this results in $r_0 = 1.5061$ Å and $k = 26.041\,\mathrm{eV/Å^2}$.

Force-based (GAP). We calculate the mean force (MF) between the two CG sites for a number of equally spaced distances on a grid. We consider two grids, one from 1.34 Å to 1.66 Å and an extended grid from 1.28 Å to 1.85 Å, with grid points spaced 0.002 Å apart. For a one-dimensional model, there is no distinction between the monomer term used in GAP and a spline-based interaction between the two sites. Hence, GAP and MSCG/FM give equivalent results. For the monomer, we will only show the GAP results.

Using the grid MF, we teach the monomer descriptor with the following hyperparameters: length scale $\theta = 0.3$ Å, energy scale $\delta = 5\,\mathrm{eV}$, 10 sparse points (using uniform selection), and a noise level of $\sigma_f = 1\,\mathrm{meV/Å}$ for force observations.

For comparison, we also teach the GAP from 10^5 samples of instantaneous collective forces (ICF). With the exception of the noise level, all hyperparameters relate to the true PMF and do not need adjustment. For the methanol monomer we can plot the distribution of CG bond forces as a function of CG bond length, shown in Fig. 5.3 on the following page. This allows us to *measure* the noise in the ICF, given here by the standard deviation of the residuals. Hence, this directs us to use $\sigma_f = 0.25\,\mathrm{eV/Å}$.

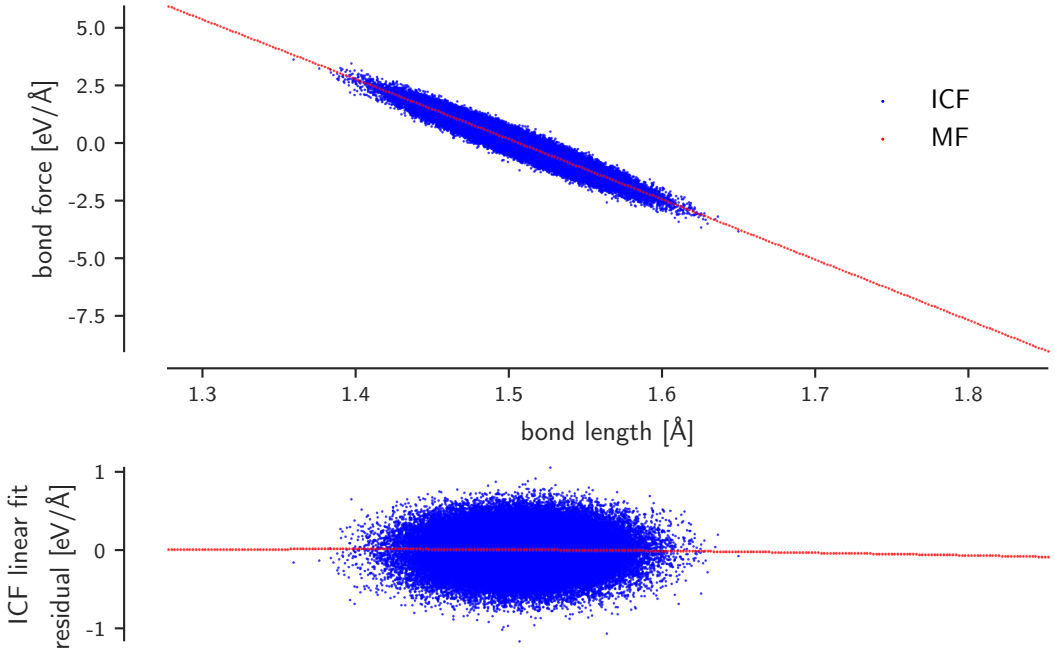

Figure 5.3: Methanol monomer training data. *Top:* CG bond force as a function of bond length for instantaneous collective forces (ICF, blue) and mean forces (MF, red). *Bottom:* residuals of ICF and MF from a linear fit to the ICF. Note that the ICF have a much higher variance than the MF, but that their average reproduces the MF. The standard deviation of the ICF residuals is 0.23 eV/Å.

In Fig. 5.4 we show the CG forces resulting from different approaches to obtaining the bond potential, and their deviation from the MF reference values. These results show clearly that the ICF-trained GAP only provides accurate results where there is sufficient data coverage. There seems to be a slight systematic difference to the MF.

The results for ICF- and MF-taught GAP demonstrate that the GAP-predicted variance is a useful measure for checking adequate data coverage, but that its absolute value needs to be treated with caution. As the ICF-trained GAP is based on a much larger number of training points, the predicted variance is smaller, even though the force prediction is generally more accurate when training from mean forces. Note that in the limit of high data coverage, the predicted variance goes towards the regularisation constant, in this case 10^{-6}.

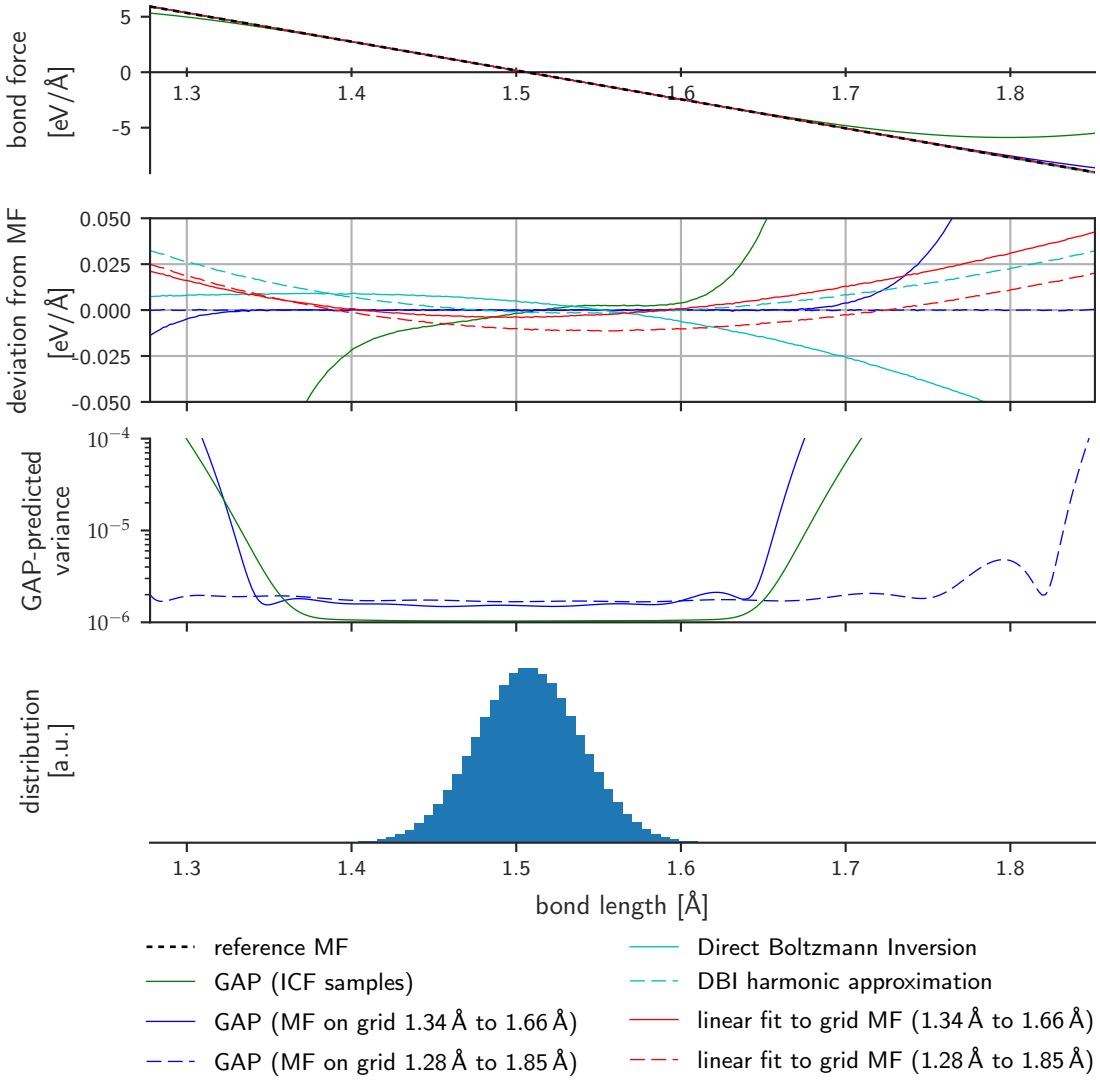
bond force [eV/Å]
deviation from MF [eV/Å]
GAP-predicted variance
distribution [a.u.]
bond length [Å]
reference MF
GAP (ICF samples)
GAP (MF on grid 1.34 Å to 1.66 Å)
GAP (MF on grid 1.28 Å to 1.85 Å)
Direct Boltzmann Inversion
DBI harmonic approximation
linear fit to grid MF (1.34 Å to 1.66 Å)
linear fit to grid MF (1.28 Å to 1.85 Å)

Figure 5.4: Methanol monomer CG potentials. From top to bottom (as a function of the CG bond length): CG bond force predictions; prediction error, as measured by the deviation from the reference mean forces (MF); GAP-predicted variance; distribution of the CG bond lengths in the all-atom simulation.
We compare the following potentials with the reference MF: GAP taught from MF (on grids with two different extents), GAP taught from instantaneous collective forces (ICF), linear fit to grid MF, and Direct Boltzmann Inversion (DBI) of the bond length distribution as well as its harmonic approximation.
The GAP-predicted variance clearly indicates the range covered by the training data. Note that its absolute value is meaningless: it is smaller for the ICF-taught GAP simply because this involved many more training points. Beyond the extent of the training data, it can be seen that the GAP regresses to the prior mean function (that is, zero).

The monomer in other environments

While the methanol monomer is not particularly interesting in itself, it provides a useful basis for generating GAP potentials for larger systems.

In our work, we found that the monomer (bonded) forces are significantly larger than the (non-bonded) forces due to two- or three-body effects. The energy scale of the monomer PMF (based on the range of GAP predictions for the monomer term) is of the same order of magnitude as that of the dimer (based on the well depth of the centre-of-mass PMF). However, the length scale associated with the monomer, corresponding to fluctuations in the CG bond length, is much smaller [on the order of 0.1 Å] than the typical length scale for dimer interactions [on the order of 1 Å]. Hence the magnitude of the forces produced by the monomer term is much larger, and can potentially mask the signal of dimer and trimer interactions.* Hence, a good monomer term is essential for learning accurate dimer and trimer potentials.

The monomer bond length distribution is a useful consistency check in CG simulations to help us determine whether something has gone wrong. Obtaining an accurate bond length distribution is a necessary (though not sufficient) condition for an accurate simulation.

The monomer term does not change significantly between systems, whether in vacuum, small clusters or bulk, and so there is little change in the distribution of the CG bond length between different methanol systems. In all cases the bond length has an average of approximately 1.507 Å and a standard deviation of 0.0315 Å, with a small trend towards shorter bond lengths in bulk. The relative change of the average bond length between vacuum and bulk is less than 5×10^{-4}. In all cases, the harmonic approximation to Direct Boltzmann Inversion gives a force constant of approximately 26.0 eV/Å^2, with no significant difference between environments. Differences between the monomer in vacuum and the monomer term in dimer, trimer, and larger systems could be captured by the dimer or trimer descriptor. However, for systems such as the bulk it might be preferable to directly train the bulk-induced monomer together with dimer and trimer terms.

*This is true for our models of methanol and benzene. This may not be the case in other systems, for example when considering highly coarse-grained models (such as in Ref. 88), where the scale of distance fluctuations within a monomer is comparable to the intermolecular distances.

5.2.2. Dimer

Atomistic structure

Two methanol molecules in vacuum form a bound dimer state through the hydrogen bond between the two hydroxy groups. In the CG model, the hydrogen bond is best described by the vector connecting the two OH sites. In the CG view, this leads to a sharp peak in the OH–OH distance distributions that extends to about 3.4 Å, with its maximum at 2.7 Å.

Due to random fluctuations, the dimer can dissociate. In a true vacuum simulation, the dissociated monomers would drift apart with vanishing probability of encountering each other again. The state with two free-floating monomers is of no further interest. We could place the dimer inside a periodic box, so that the two molecules will eventually encounter each other and be able to form a bound dimer again. However, this is a finite-size effect; in an arbitrarily large box, the probability of re-forming the dimer state goes to zero. Instead, to improve the sampling of the bound dimer state, we impose a half-open harmonic restraining potential on the distance r_{COM} between the centres of mass of the molecules,

$$U_{\mathrm{rst}}(r_{\mathrm{COM}}) = \begin{cases} 0 & r_{\mathrm{COM}} \leq r_{\mathrm{rst}}, \\ \frac{1}{2} k_{\mathrm{rst}} (r_{\mathrm{COM}} - r_{\mathrm{rst}})^2 & r_{\mathrm{COM}} \geq r_{\mathrm{rst}}. \end{cases} \tag{5.2.1}$$

We let the restraint begin at $r_{\mathrm{rst}} = 10\,\text{Å}$, and we set the force constant of the restraining potential to $k_{\mathrm{rst}} = 50\,\mathrm{kcal/mol/Å^2}$. The same restraining potential is applied to all-atom and CG simulations.

In the two-site CG model, the methanol dimer has six degrees of freedom. These can be completely described by the two bonded and four non-bonded site–site distances as in the dimer descriptor. To compare the different models, we will show distributions of collective variables of the CG model that are easier to understand, such as the angle between a CG bond and the OH–OH separation vector and the dot product between the two normalised bond directions.

CG potentials

Here we compare the GAP approach based on molecular terms with the MSCG/FM approach based on pair potentials. Both approaches are based on forces, and we use a training set containing 11400 dimer configurations that were randomly sampled from

the atomistic trajectory and a further 4894 dimer configurations sampled at short range (with a dimer COM separation less than 3.3 Å).

GAP. The methanol dimer is completely described by the dimer descriptor, which for this system contains six distances (two bonded and four non-bonded).

As discussed in Section 4.2.4, it is advantageous to separate monomer and dimer contributions by explicitly including a monomer descriptor, as the largest contribution to the mean forces comes from the monomer bond forces. This can be seen in Fig. 5.5, where we compare mean forces in the dimer system with the force prediction of the monomer GAP discussed in the previous section.

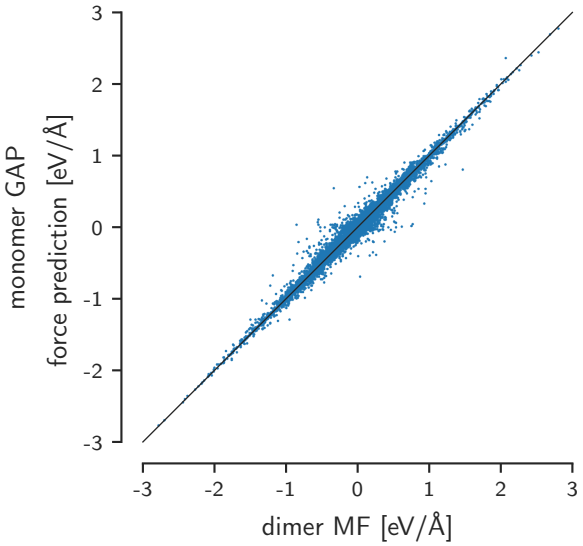

Figure 5.5: Comparison between monomer and dimer forces. This graph shows that a large contribution to the mean forces in the methanol dimer is already explained by the force predictions given by the monomer term. Two-body contributions to the dimer mean forces show up as deviations from the diagonal line.

In principle, GAP can automatically split the contributions associated with different descriptors (and we will make use of this for the bulk systems). Here, we simply use the monomer GAP discussed in the previous section as a core potential. On top of this core potential, we teach the dimer descriptor to capture the two-body contributions to the mean forces that are not described by the monomer term.

As discussed in Section 4.3.3, at short range the interaction between atoms and, hence, between CG sites becomes highly repulsive, which may lead to insufficient

sampling and unphysical behaviour of the CG interactions. In our approach, we handle this by introducing a short-range repulsive interaction between pairs of CG sites. We generate the repulsive pair interaction from mean force calculations on a grid, where the different combinations of CG sites (CH_3–CH_3, CH_3–OH and OH–OH) are brought together. This is then combined with the monomer GAP discussed previously. The resulting core potentials are shown in Fig. 5.6.

As discussed in Section 4.3.2, we need to take into account the change in length scale between long-range and short-range interactions. To this end we apply the transfer function Eq. (4.3.2) to all non-bonded distances in the dimer descriptor. The hyperparameters for the dimer-descriptor GAP are listed in Table A.1 in the Appendix.

MSCG/FM pair potentials. For comparison, we also apply the spline-based MSCG/FM approach to the methanol dimer. This is equivalent to a GAP using pairwise distance descriptors. To describe the interactions between two-site methanol monomers in terms of pair potentials between sites, we need one spline interaction for the CG bond and three non-bonded splines. We obtained the splines for these interactions using the MSCG/FM code provided by the Voth group. As a consistency check, we also taught a GAP with the corresponding four pairwise distance descriptors; this leads to the same potentials. The resulting non-bonded pair potentials are shown in Fig. 5.6. (The bonded pair potential is equivalent to the monomer potential and is not shown again.)

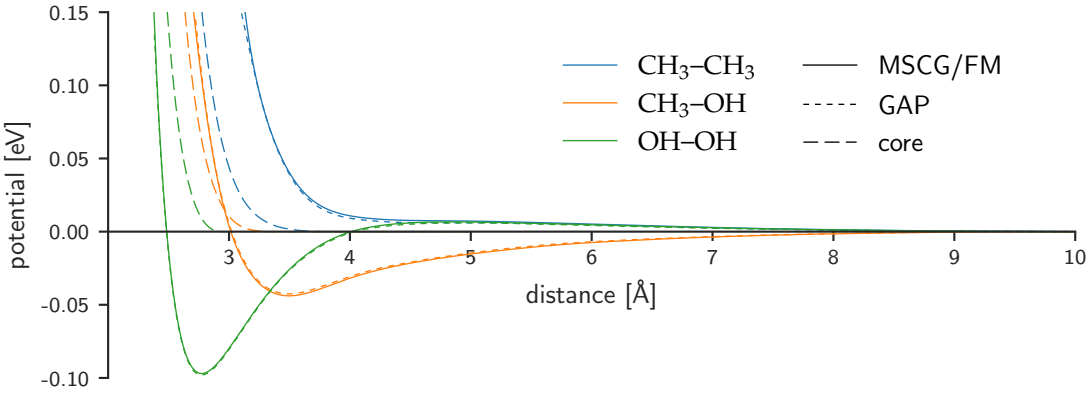

Figure 5.6: Methanol dimer pair potentials. The non-bonded interactions are described by three pair potentials for the CH_3–CH_3, CH_3–OH, and OH–OH interactions. Here we show the spline potentials obtained in the MSCG/FM approach (solid lines) and, for comparison, the potentials obtained in the GAP approach using one-dimensional pairwise distance descriptors (dotted lines). This shows the equivalence to the MSCG/FM approach. The highly attractive hydrogen bond between two OH sites results in the much deeper well depth for the OH–OH pair potential. We also show the short-ranged core potentials used in the training of the dimer-descriptor GAP (dashed lines).

The MSCG/FM approach is less reliant on core potentials because the splines do not have the Gaussian process's tendency to drive the energy back to zero far away from the data. Moreover, due to the much lower dimensionality of pairwise splines compared to the molecular descriptors, they are much easier to saturate with data. In all our tests, we did not run into issues due to insufficient sampling for the MSCG/FM splines. We note, though, that in practice the MSCG/FM approach often uses a "core potential" to ensure short-range repulsion as well (see, for example, Ref. 88).

Comparison of potentials

While it is not feasible to visualise predictions of the potentials in the full six-dimensional space that describes the methanol CG dimer, we can display forces along cutlines through this space. Here, we consider fixed relative orientations of the methanol molecules and vary their COM separation. In Fig. 5.7 we compare the COM force predictions of dimer GAP and pair potentials with mean force reference values. This shows that GAP captures the MF with very small error, whereas the pair potentials exhibit a visible discrepancy. In Figs. 5.8 and 5.9 we show the effects of transfer function and short-range core potential on the dimer GAP predictions.

We also compare the force prediction error of the CG potentials by measuring the

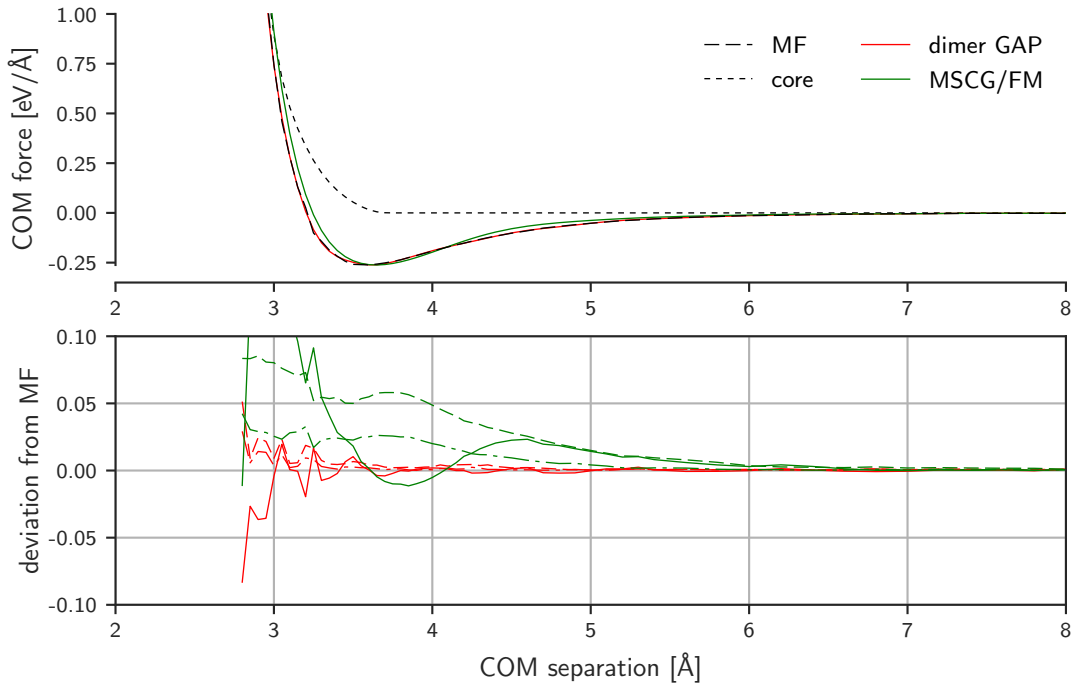

Figure 5.7: Methanol dimer: forces along a cutline (1/3) — comparison of dimer-descriptor GAP and pair potentials. *Top:* centre-of-mass force between the molecules along the COM separation, comparing CG predictions with MF. *Bottom:* deviation of CG potential predictions from MF [solid: COM force along COM separation; dash-dotted: perpendicular (shearing) component of COM force; dashed: torque-inducing force (not along a CG bond)]. This shows the core potential, how GAP captures the MF with very small error, and how the MSCG/FM pair potentials produce forces significantly different from the MF.

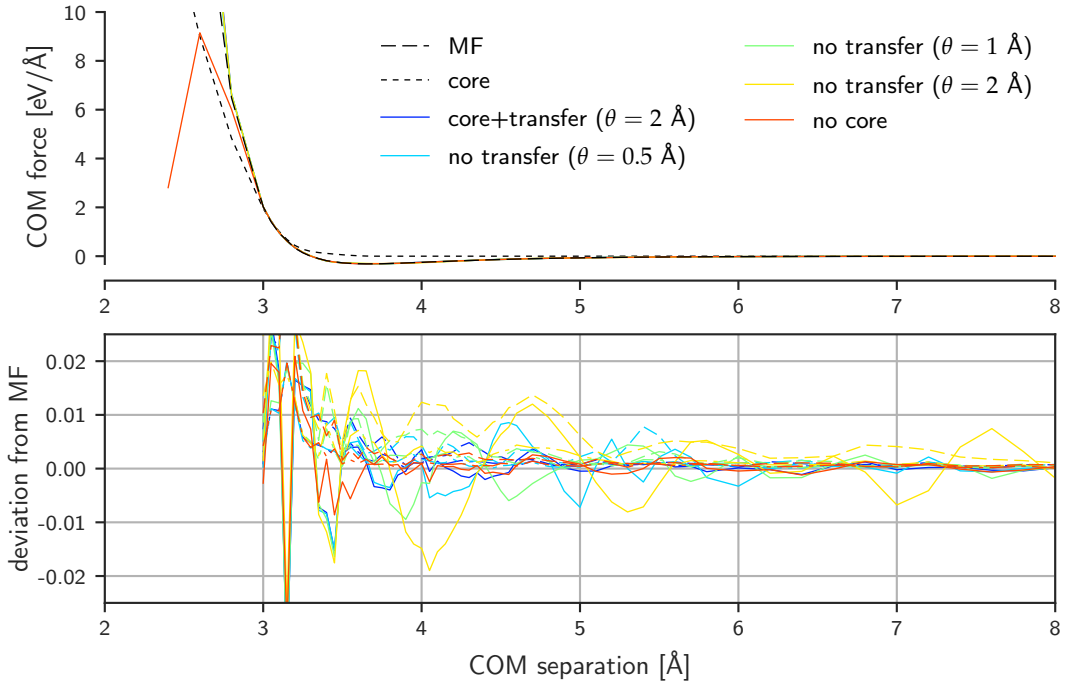

Figure 5.8: Methanol dimer: forces along a cutline (2/3) — transfer function and core potential. This shows that, without a transfer function, there is no single length scale θ that is appropriate for the whole range of COM separations which, in turn, is responsible for the resulting "ringing" that can be clearly observed along the cutline. It also shows how, without a repulsive core potential, the force prediction peaks and then turns to zero (and will eventually become negative) in the regions where there is insufficient sampling.

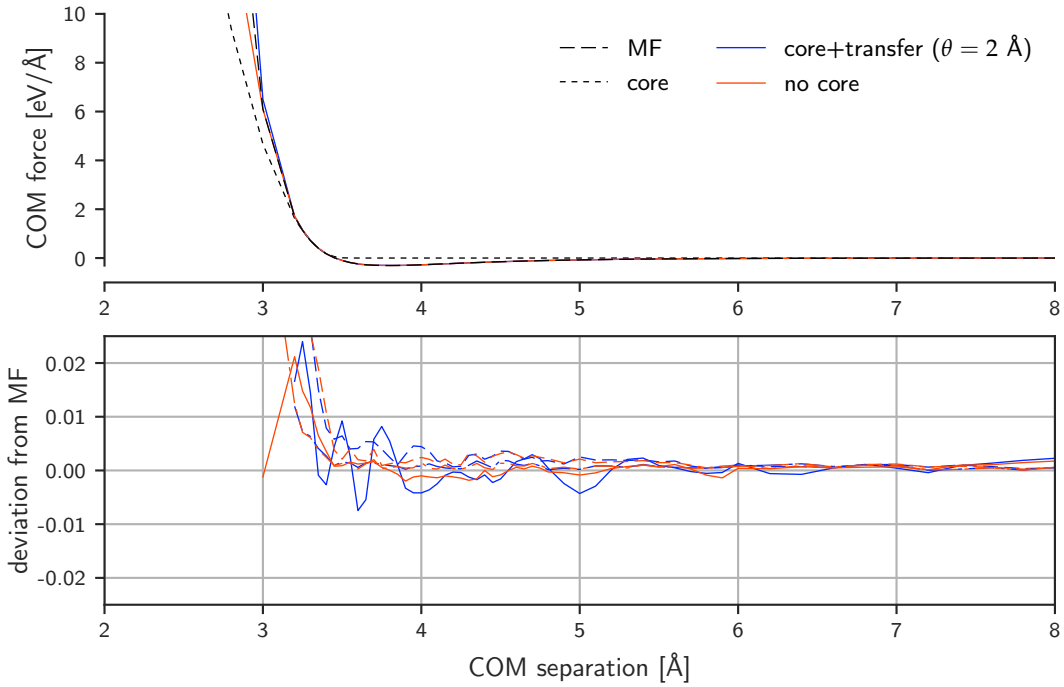

Figure 5.9: Methanol dimer: forces along a cutline (3/3) — drawback of core potential. Here we can see one drawback of using core potentials: the GAP without the core potential actually captures the dimer cutline more accurately. Though the aim of the core potential is to make the function that remains to be fitted smoother, and specifically to reduce the steep short-range repulsion, here it may have led to a somewhat more rugged function that is then slightly less well-represented by the GAP. This may be due to the repulsive OH–OH core potential extending further than the full MSCG/FM pair potential, as can be seen in Fig. 5.6. This issue will be investigated in more detail in future research.

deviation from the components of the MF reference on a separate test set of 1400 randomly sampled dimer configurations that were not included in the training. In Fig. 5.10 we show this error as a function of the two-body contribution to the force magnitude. Note that the error for the GAP predictions is distributed symmetrically around zero. In contrast, the pair potential predictions have a much larger error that also contains a systematic component. In Fig. 5.11 we show the force prediction error as a function of the dimer COM separation, compared to the two-body contribution to the force magnitude. The two-body force magnitude decreases with the COM distance. This shows that the relative error of the GAP forces is much smaller than that of the forces given by the pair potentials. The root mean squared error (RMSE) for the different potentials across the whole test set is given in Table 5.1.

Whereas errors are usually assumed to follow a Gaussian distribution, we note that in the case of the clusters the deviation between CG potential prediction and mean force reference is more closely described by a Laplace distribution: in a logarithmic histogram of the distribution of force component deviations, the tails show a linear trend. (In bulk methanol, the errors follow a Gaussian distribution as expected.)

potential	derived from	RMSE / meV/Å
monomer GAP	monomer MF	60
dimer GAP	dimer MF	5
MSCG/FM pair potentials	dimer MF	23

Table 5.1: Methanol dimer force prediction errors. Comparison of the force prediction root mean squared error (RMSE) on a separate MF test set (1400 configurations). We compare the monomer GAP, the dimer GAP, and the dimer pair potentials.

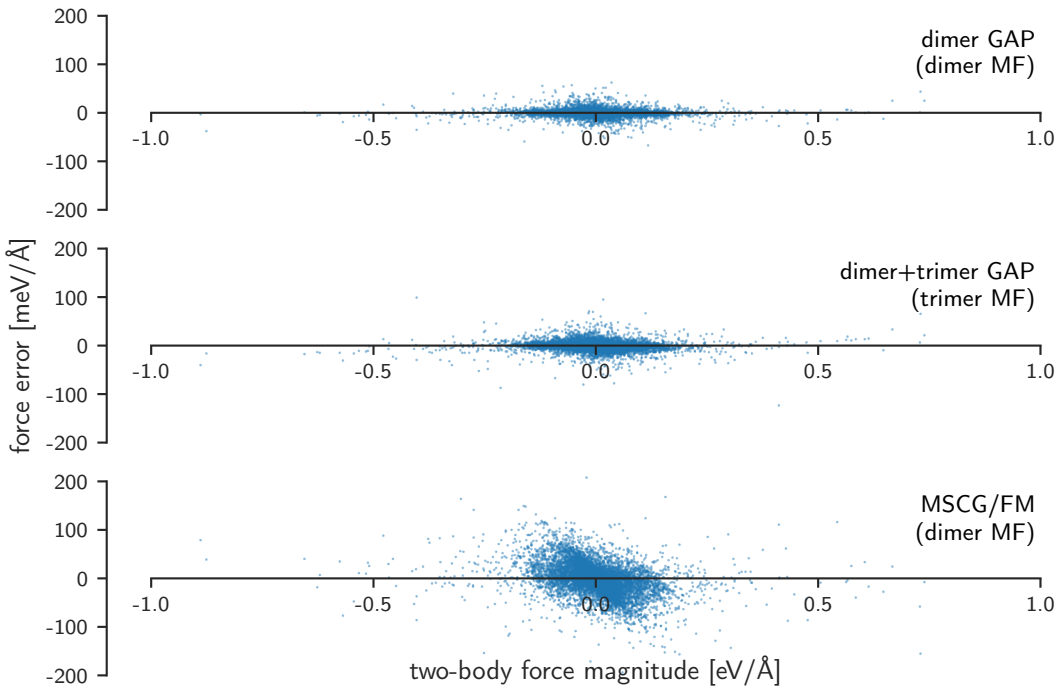

Figure 5.10: Methanol dimer: force prediction error *vs.* force magnitude. Here we compare the deviation of the force predictions of CG potentials from MF on a separate test set as a function of the two-body contribution to the force component magnitude. *Top:* dimer-descriptor GAP taught from dimer MF. *Middle:* dimer prediction of a GAP with dimer and trimer descriptors simultaneously taught from trimer MF (*cf.* Section 5.2.3). *Bottom:* MSCG/FM pair potentials from dimer MF. This shows that the largest errors do not occur for the largest forces, but rather at small force magnitude. Moreover, whereas the error in the GAP predictions is small and symmetrically distributed around zero, the pair potential predictions contain a systematic error.

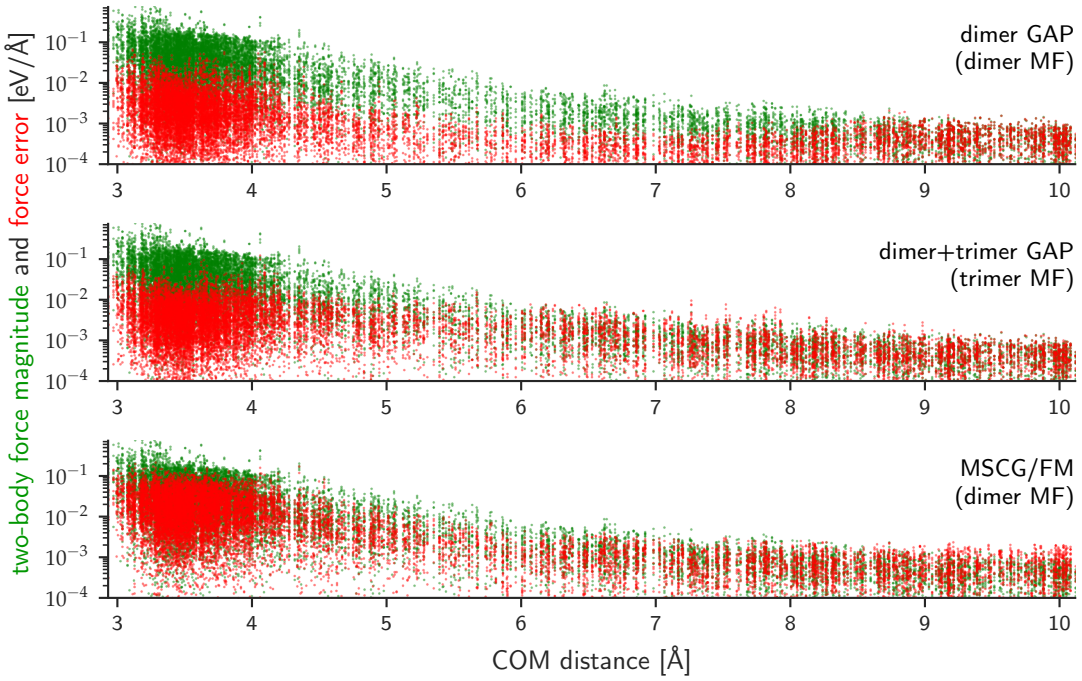

Figure 5.11: Methanol dimer: force prediction error *vs.* COM distance.
Here we show the force prediction error as a function of COM separation (red) and, for comparison, the absolute magnitude of the mean force components (green). *Top:* dimer-descriptor GAP taught from dimer MF. *Middle:* dimer prediction of a GAP with dimer and trimer descriptors simultaneously taught from trimer MF (*cf.* Section 5.2.3). *Bottom:* MSCG/FM pair potentials from dimer MF.

MD results

While it may be useful to assess the accuracy of CG potentials by comparing their predictions with a mean force test set, the real test of CG potentials is their performance in MD simulations. We now consider the results of simulations using the GAP and MSCG/FM CG potentials. We compare distributions in the CG ensembles with the corresponding distributions of the all-atom simulation.

Previous research has generally focused on the RDF for molecular centres of mass and for pairs of CG sites. Structure-based coarse-graining reproduces the site–site RDFs by construction. Inadequacies of these approaches only show up when considering more complex distributions and correlations between them [135]. Hence, here and in the following we present a wide range of observables that can be obtained from the CG coordinates, including two-dimensional correlation plots.

In Fig. 5.12 we show the distributions of COM separation and site–site distances for all-atom, GAP-CG and MSCG/FM simulations. In Fig. 5.13 we show distributions of the angle α between a CG bond and the hydrogen bond and of the dot product $n_1 \cdot n_2$ between the normalised directions of the two CG bonds. These are captured to a very high accuracy by the dimer GAP, whereas the MSCG/FM results show significant discrepancies. Figure 5.14 further demonstrates that GAP is able to reproduce the correlation of α and $n_1 \cdot n_2$, which is not the case for pair potentials.

This demonstrates that the dimer GAP can perfectly reproduce many kinds of distributions beyond simple RDFs. Indeed, any discrepancies are due to insufficient sampling. While the pair potentials provide a reasonable description of the methanol dimer, they show systematic deviation from the all-atom distributions.

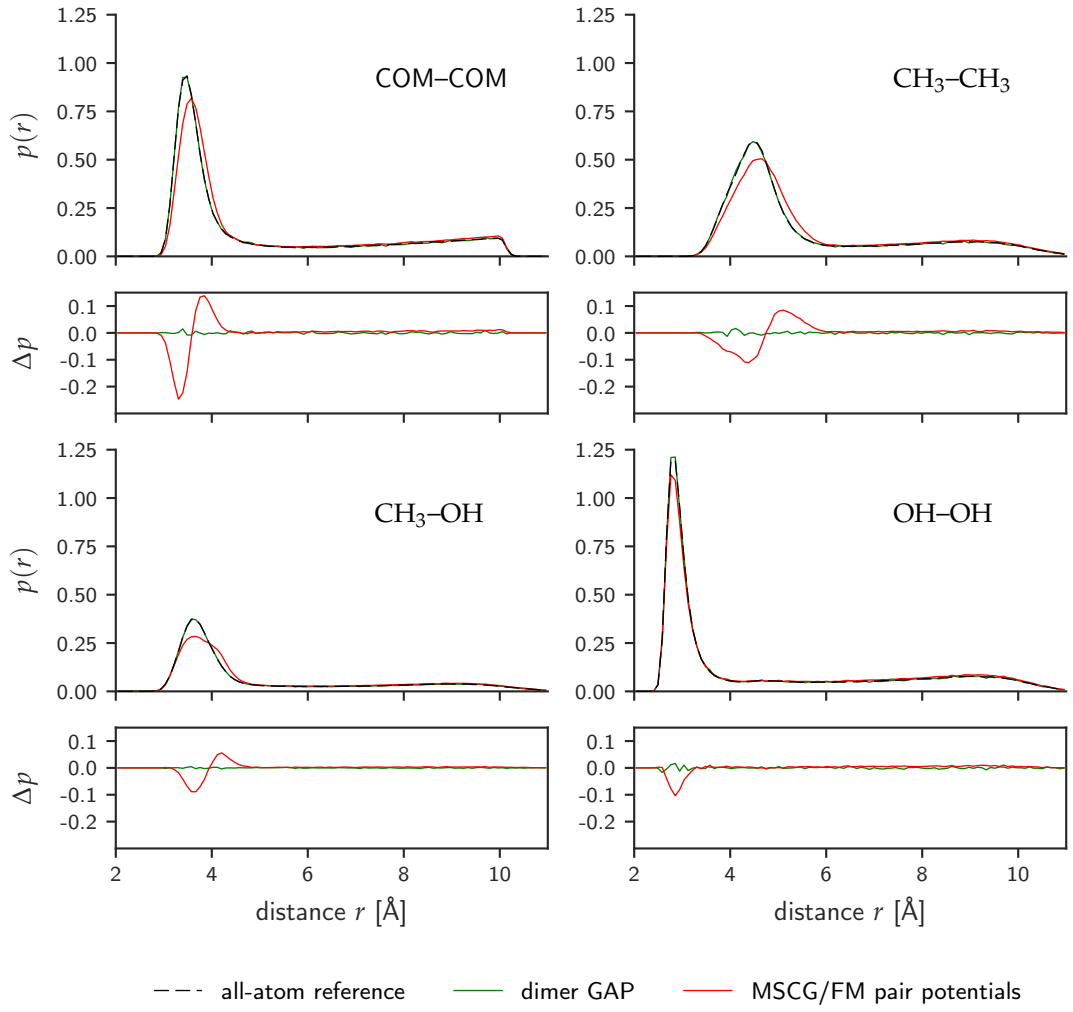

Figure 5.12: Methanol dimer: distance distributions. Normalised distance distributions $p(r)$ for COM distance (top left), CH_3–CH_3 distance (top right), CH_3–OH distance (bottom left) and OH–OH distance (bottom right). We compare MD results for all-atom model, dimer GAP CG model, and MSCG/FM CG model. Δp is the deviation from the reference distributions in the all-atom simulation.

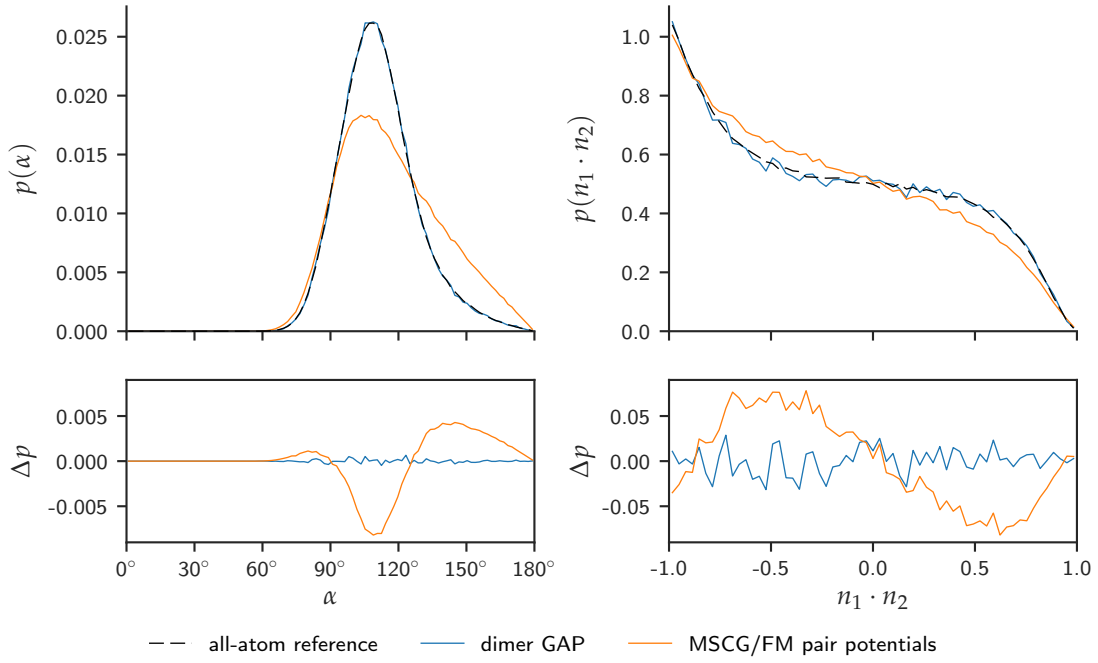

Figure 5.13: Methanol dimer: orientational distributions. *Left:* distribution of the angle α between CG bond and hydrogen bond/OH–OH axis (CH3–OH$\cdots$OH angle). The peak in $p(\alpha)$ coincides with the equilibrium value of the C–O–H bond angle in the all-atom model, 108.5°. *Right:* distribution of the dot product $n_1 \cdot n_2$ between CG bond directions. If the two bond directions were uncorrelated, the dot product $n_1 \cdot n_2$ would be distributed uniformly. Only bound dimer configurations with a OH–OH distance less than 3.4 Å were included in these distributions.

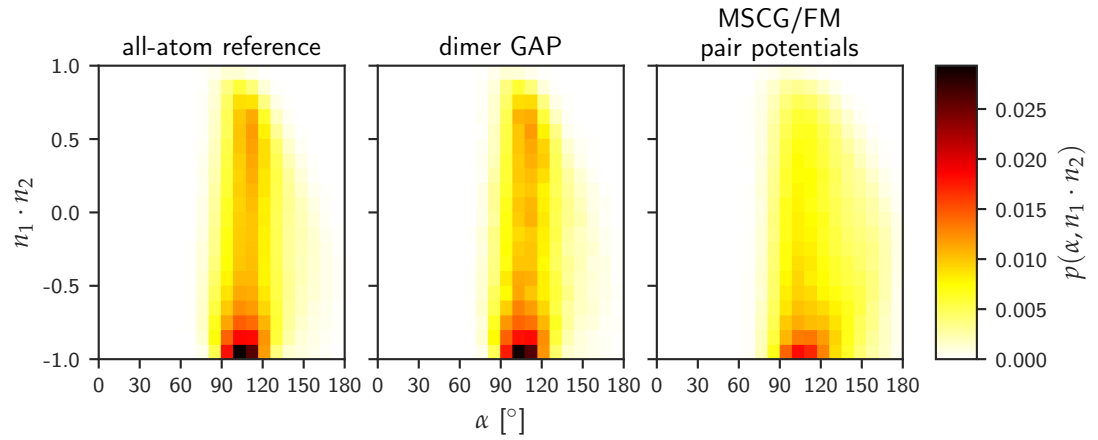

Figure 5.14: 2D correlation between α and $n_1 \cdot n_2$ (*cf.* Fig. 5.13).

5.2.3. Trimer

Atomistic structure

The methanol trimer exhibits two distinct stable configurations.

On the one hand, there are configurations in which there are three hydrogen bonds that form an equilateral triangle, with all angles between OH sites approximately 60°. The molecules tend to form a planar configuration in which the bond vectors are pointing towards the centre of the triangle.

On the other hand, there are configurations with two hydrogen bonds directed towards a central molecule, forming an isosceles triangle. There, the angle between the hydrogen bonds varies between approximately 90° and 120°. Hence, the other two angles in the triangle of OH sites generally lie in the range 30° to 45°.

These two configurations can be clearly distinguished in plots of the angular distribution function. To analyse orientations of the bonds and differences between the two types of structure, it makes sense to consider these configurations separately. We define the following masks:

"Bound configuration": at least two OH–OH distances are less than 3.4 Å, indicating a hydrogen bond.

"Equilateral configuration": bound configuration where all angles within the triangle of OH sites are larger than 55°; this implies that any angle can be at most 70°.

"Isosceles configuration": bound configuration where two of the angles differ by at most 5°, and the third angle is larger than 70°.

We show the distribution of angles within bound configurations in Fig. 5.15, where we also highlight those configurations that we count as "equilateral" or "isosceles".

As in the case of the dimer, we employ half-open distance restraints between the centres of mass of the molecules to keep the cluster from dissociating. Here we apply restraints between each pair of molecules, with the restraint setting in at 15 Å. The force constant is 50 kcal/mol/Å^2.

CG potentials

In the all-atom model used for this work, adding more molecules to the system does not change the form of each individual interaction. If this were true for the CG model, we could simply take the potentials obtained from dimer MF as described in the previous

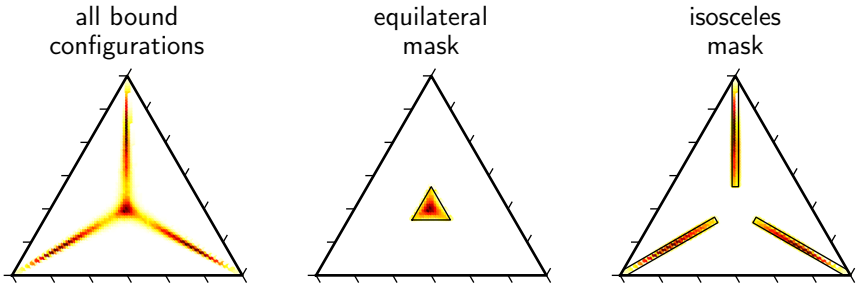

Figure 5.15: Methanol trimer: distribution of angles between OH sites and angle masks in all-atom model. *Left:* ternary plot of the angles in the triangle formed by OH sites. Each axis represents one of the angles from 0° to 180°. (Only bound trimer configurations with at least two OH–OH distances less than 3.4 Å are included.) *Middle:* the mask used to capture equilateral configurations. *Right:* the mask used to capture isosceles configurations.

section and apply them to systems containing any number of molecules. This would also imply that we would obtain the same interactions whether teaching the dimer GAP from dimer MF or trimer MF. Unfortunately, this is not the case, as we will see in the deviation of the predicted forces from the all-atom mean forces along a dimer cutline for a dimer-descriptor GAP trained on *trimer* MF, shown in Fig. 5.16. In this case the dimer term will acquire an effective contribution due to the actual trimer interactions.

GAP. We compare three different GAPs. One is the dimer GAP trained on dimer MF that was discussed in the previous section. Furthermore, we consider two GAPs taught from trimer MF, one using only the dimer descriptor and one using the trimer descriptor (using the dimer GAP trained on dimer MF as a core potential). We will also discuss a simultaneous teaching of dimer and trimer descriptor from trimer MF.

For teaching a dimer GAP from trimer MF we use the same hyperparameters as for the methanol dimer GAP (see Table A.1 in the Appendix). As the dimer descriptor does not capture the full dimensionality of the trimer system and cannot exactly represent the trimer PMF, there is an additional, systematic error beyond just the noise in mean forces. Hence, in principle, the noise parameter should be increased, corresponding to an increased amount of regularisation. In this case, increasing the noise level hyperparameter turned out not to have any noticeable effect in practice.

The trimer descriptor of the two-site methanol contains 15 distances (three bonded and 12 non-bonded). For the trimer-descriptor GAP used in MD simulations, we use

the dimer GAP that was presented in the previous section as the core potential. The hyperparameters for the trimer descriptor are given in Table A.2 in the Appendix. It is also possible to teach the dimer and trimer descriptors simultaneously from trimer MF. For this we use the hyperparameters for dimer and trimer descriptor given in Tables A.1 and A.2.

In Fig. 5.16 we compare the predicted forces along a dimer cutline for the different GAP potentials with the dimer MF reference values. Teaching only the dimer descriptor from trimer MF results in a significantly different curve for the forces along the cutline. This shows that the dimer term picks up an effective contribution due to many-body effects. In the following we will see that neither of the potentials using only the dimer descriptor is able to reproduce the atomistic structure correctly. In our two-site methanol trimer, there are genuine three-body effects that are not captured by the dimer descriptor alone; they can only be described by considering the molecular three-body interactions. Teaching dimer and trimer descriptors simultaneously allows the dimer descriptor to recover the actual dimer cutline, while the trimer descriptor captures all three-body effects. This supports our decision to separate the PMF into molecular monomer-based terms.

MSCG/FM pair potentials. The methanol trimer system contains the same types of pairwise interactions as the dimer. However, we do not assume that the dimer pair potentials accurately describe the trimer system. Hence we re-calculate the splines from trimer MF. The resulting potentials are shown in Fig. A.1 in the Appendix.

Structure-based pair potentials (IBI). For the methanol trimer, we include results for structure-based pair potentials that reproduce the site–site distance distributions. As mentioned in Chapter 3, different approaches such as IBI and IMC would arrive at the same pair potentials; as we consider comparatively small systems, for the present work, IBI is our method of choice.

For technical reasons, it was impractical to carry out IBI directly on the trimer in vacuum with restraining potential.* Hence, we placed the trimer inside a box with side length $L = 26\,\text{Å}$ and periodic boundary conditions (PBC); this simulation produced sufficient data to obtain converged RDFs for pairs of CG sites.

In a periodic box, the long-range electrostatic interactions of the partial charges lead to interactions with their mirror images. This results in forces on the atoms (and

*The VOTCA package used to carry out the Iterative Boltzmann Inversion was not able to calculate distance distributions for the open boundary conditions of the vacuum simulations.

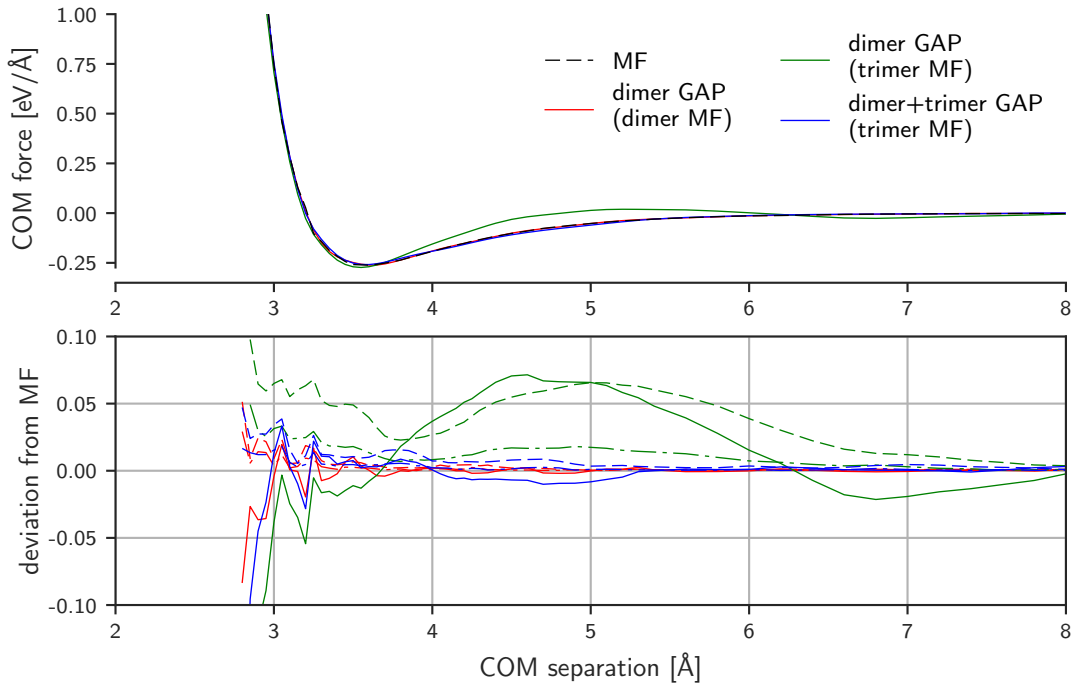

Figure 5.16: Methanol trimer GAP: force prediction along dimer cutline. Training only a dimer-descriptor GAP from trimer MF results in a systematic deviation of the dimer force prediction from the MF reference. In contrast, training both dimer and trimer descriptors simultaneously recovers the dimer forces along the cutline very well. For comparison, we include the dimer-descriptor GAP trained from dimer MF.

hence, mean forces on the CG sites) that are slightly different from those in a vacuum simulation with open boundary conditions (and otherwise equivalent settings). Hence, the GAP taught from vacuum MF does not, and is not expected to, exactly reproduce the all-atom distance distributions in a periodic box. To overcome this problem, we could have recalculated MF in the atomistic periodic box simulation and trained a new GAP from these forces. However, here we will instead present results for GAP and MSCG/FM (compared to the all-atom simulation with restraints) separately from results for IBI (compared to the all-atom simulation with PBC). In any case, this is only a limitation with respect to overall distributions of distance variables. When only considering configurations in which all three pair distances are within a certain cutoff ("bound states"), there is no difference between the distributions from simulations carried out in vacuum using distance restraints and those from simulations in a box with periodic boundary conditions.

The CG bond potential was obtained by direct Boltzmann Inversion using the harmonic approximation. The non-bonded pair potentials were refined iteratively using the VOTCA package [94] (based on the direct Boltzmann Inversion as a starting point). The three pair potentials were defined on a grid up to a cutoff of 10 Å with spacing 0.1 Å. To improve the stability of the IBI algorithm, potential updates were smoothed twice and scaled by a factor of 0.2. In each step, only one of the three non-bonded potentials was updated.

The resulting pair potentials accurately reproduce the site–site distance distributions of the all-atom model, as can be seen in Fig. 5.18. There is a minor discrepancy for the OH–OH distance distribution, which is due the sharpness of the peak corresponding to the hydrogen bond. To capture this peak accurately would require a much finer grid, which would in turn require much longer simulation runs to obtain converged distributions, making the calculations more difficult.

Comparison of potentials

We show the force prediction errors of the potentials on a separate trimer MF test set in Table 5.2. It can be clearly seen that the trimer-descriptor GAP trained on trimer MF has a significantly smaller error in the predictions (as compared to reference MF values) than the other potentials. However, while RMSE values can give a first indication as to the accuracy of a potential, the true test for the quality of a CG potential is its performance in MD simulations.

potential	derived from	RMSE / meV/Å		
		total force	dimer part	trimer part
monomer GAP	monomer MF	97	91	42
dimer GAP	dimer MF	43	10	42
dimer GAP	trimer MF	37	20	42
dimer+trimer GAP		24	12	25
trimer GAP		26	10	25
MSCG/FM pair potentials	dimer MF	52	37	42
	trimer MF	50	38	42
IBI pair potentials	trimer RDF	57	44	42

Table 5.2: Methanol trimer force prediction errors. Comparison of the force prediction root mean squared error (RMSE) on a separate mean force test set for different potentials. We also compare the RMSE of those contributions to the total force that are only due to two-body interactions ("dimer part") and of those that are only due to three-body interactions ("trimer part").
Note that the simultaneous dimer+trimer GAP training results in a slightly smaller RMSE for the total force prediction than the trimer-descriptor GAP that uses the methanol dimer GAP as a core potential. On the other hand, the dimer+trimer GAP shows a slightly larger error for the dimer contributions to the total forces. This suggests that the split of the interactions between dimer and trimer descriptor could be somewhat improved still; this might also reduce the small discrepancy between the forces along cutlines that can be seen in Fig. 5.16.

MD results

The simulation with the IBI pair potentials was carried out in periodic boundary conditions; for comparison, we also carried out a simulation in periodic boundary conditions with the trimer GAP trained on vacuum trimer MF. All other simulations were performed using the distance restraints discussed on page 106. We show the resulting distance distributions (COM, CH_3–CH_3, CH_3–OH, OH–OH) for vacuum simulations using the distance restraints in Fig. 5.17 and for periodic boundary conditions in Fig. 5.18. Though the IBI pair potentials reproduce the site–site distance distributions by construction, they show a small but systematic deviation in the COM distance distribution. In contrast, the trimer-descriptor GAP reproduces both site–site *and* COM distance distribution of the all-atom model on whose forces it was trained. All other CG potentials show some deviation from the all-atom results. Note that there is a clear bound state for all the potentials.

As we have discussed previously, obtaining the correct distance distributions does not imply that other distributions are also captured accurately. This can be seen in Fig. 5.19, which shows ternary plots of the distributions of the angles in the triangle formed by the three OH sites. We show a slice through this distribution in Fig. 5.20; this highlights that the trimer-descriptor GAP reproduces the all-atom distribution very accurately, whereas all other CG potentials do not represent the isosceles configuration correctly.

Moreover, even though the IBI pair potential captures equilateral configurations, it does not reproduce the bond orientations correctly. This can be seen in Fig. 5.21, which shows the correlation between relative orientation and OH–OH distance for pairs of molecules separately for configurations classified as “equilateral” or “isosceles”. The trimer GAP is the only CG potential that accurately reproduces the distributions of the CG variables in the all-atom model. Note that even though the dimer GAP trained from trimer MF does not describe the isosceles configurations correctly, it reproduces the pair orientations in equilateral configurations much better than the dimer GAP from dimer MF or the pair potentials.

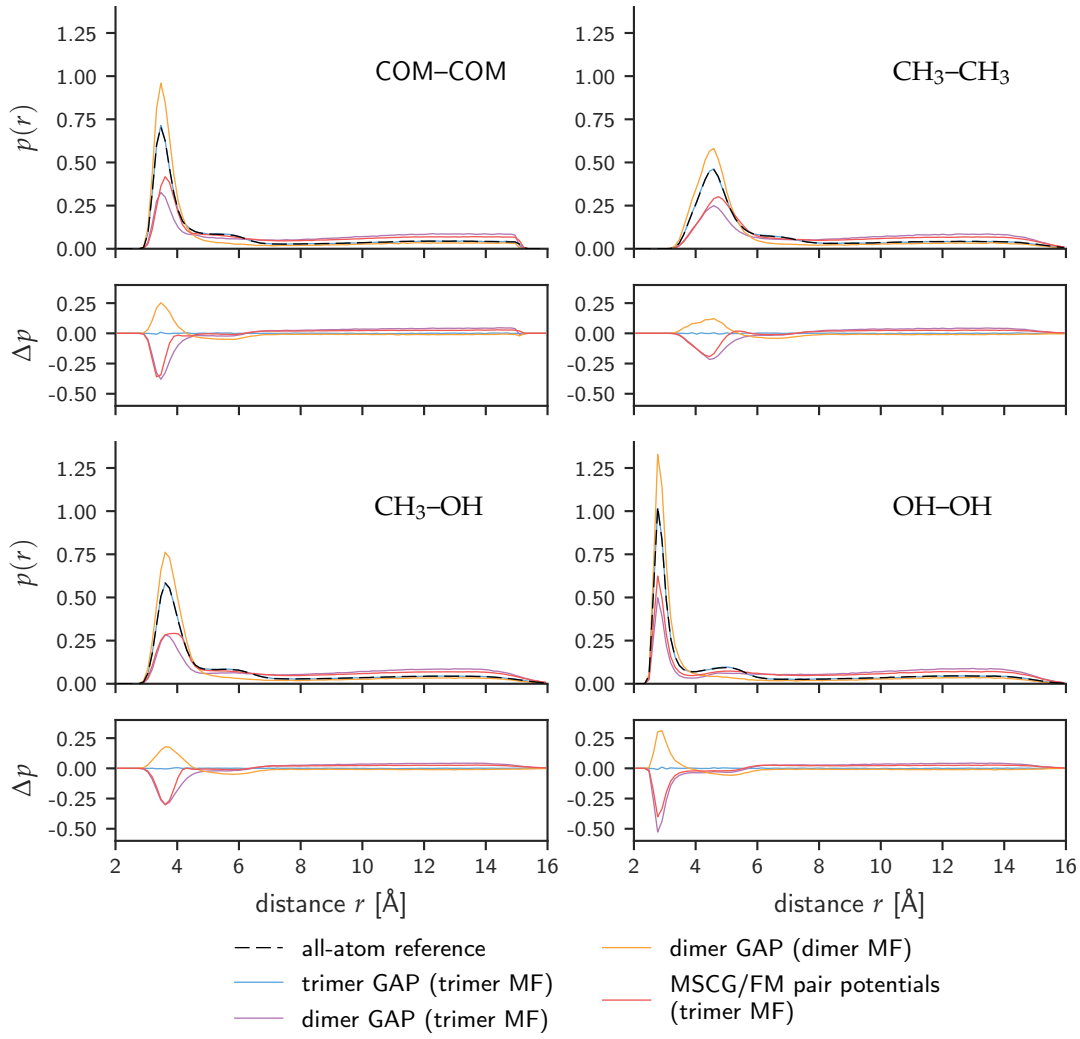

Figure 5.17: Methanol trimer: distance distributions (for simulations carried out with distance restraints). COM and site–site distance distributions $p(r)$, normalised over the range 2 Å to 16 Å, and deviation $\Delta p(r)$ from the all-atom result. Note that the distance distributions for simulations with the trimer-descriptor GAP are indistinguishable from the all-atom distributions.

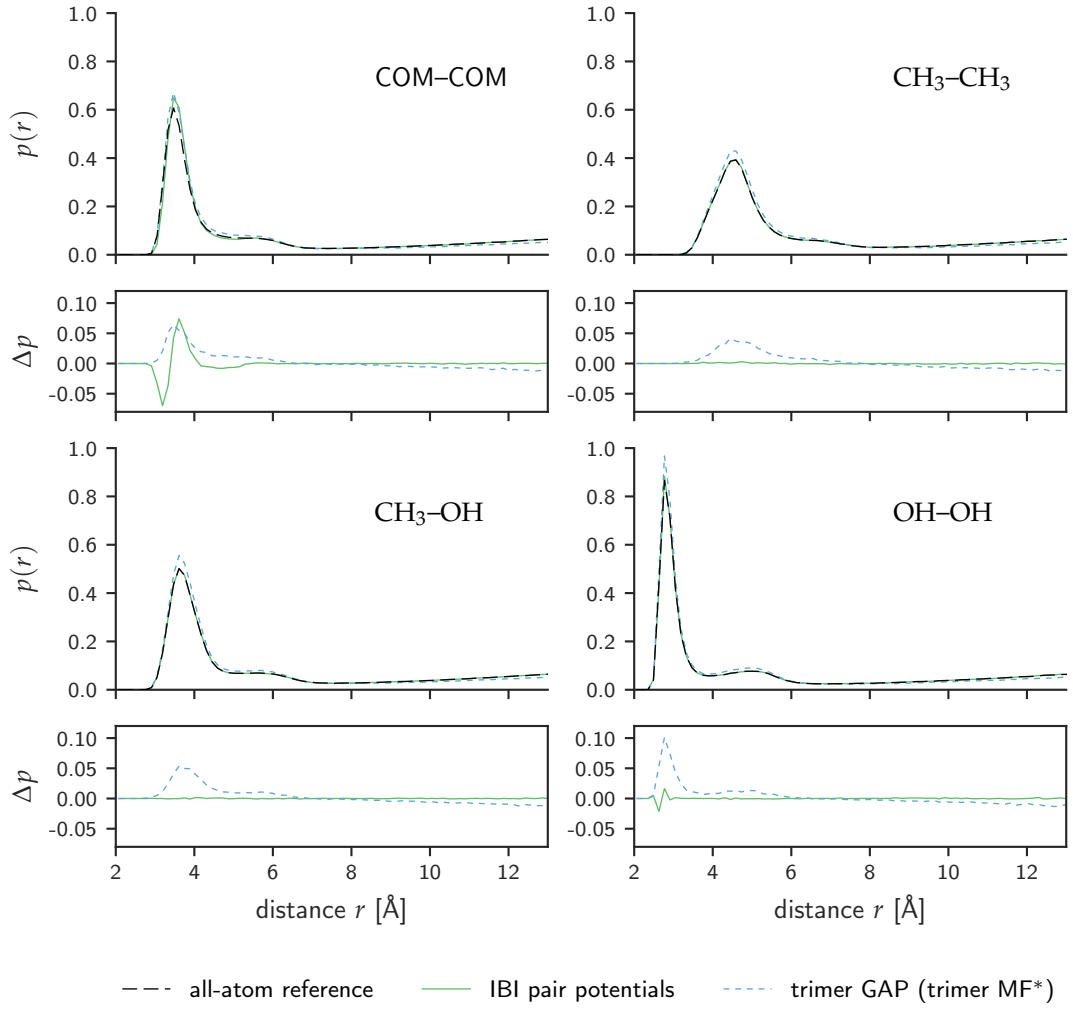

Figure 5.18: Methanol trimer: distance distributions (for simulations carried out with periodic boundary conditions). COM and site–site distance distributions $p(r)$, normalised over the range 2 Å to 16 Å, and deviation $\Delta p(r)$ from the all-atom result. Note that the IBI pair potentials exactly reproduce the site–site distance distributions (this is by construction), but nonetheless exhibit a systematic discrepancy in the COM distance distribution.
*The trimer GAP was trained on trimer MF obtained from simulations with open boundary conditions and, for this reason, does not exactly reproduce the distance distributions for periodic boundary conditions.

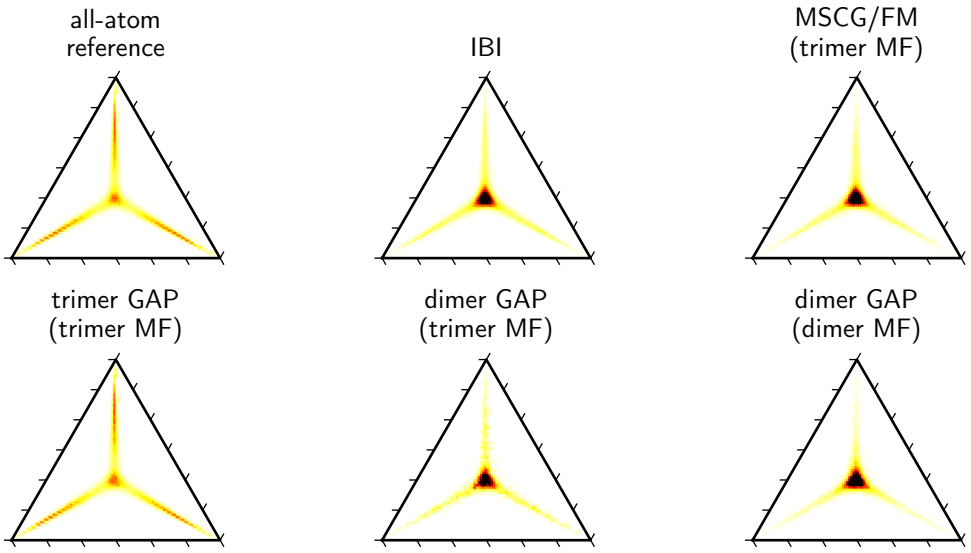

Figure 5.19: Methanol trimer: distribution of angles between OH sites. Only bound configurations (at least two OH–OH distances less than 3.4 Å) are included. The color scale is equivalent for all plots. The scale has been capped such that the secondary peaks for all-atom and trimer GAP simulations are visible; this means that the central peaks of the other plots saturate the scale.

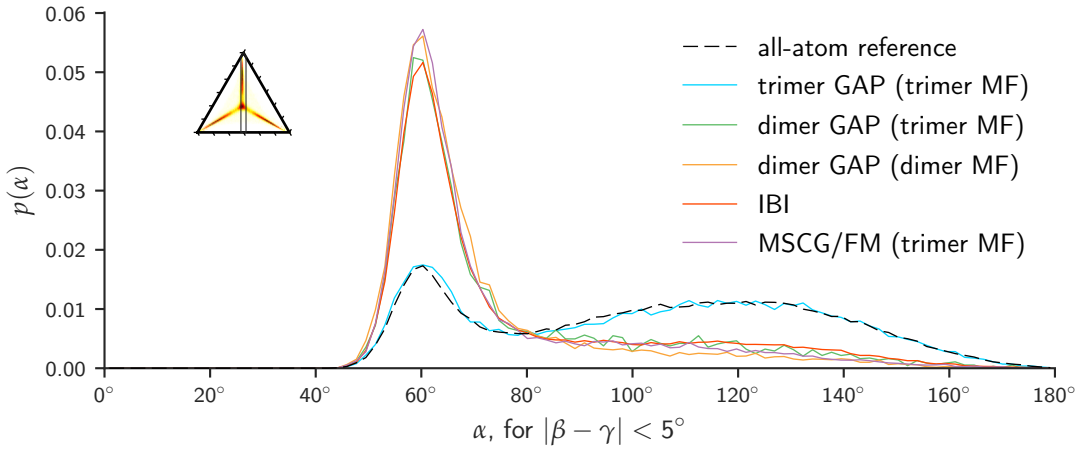

Figure 5.20: Methanol trimer: slice through the angle distributions shown in Fig. 5.19. We show the distribution $p(\alpha)$ for one angle, where the other two angles β, γ are approximately equal ($|\beta - \gamma| < 5°$). The peak at 60° corresponds to the equilateral configurations. The second, broader peak around 120° corresponds to isosceles configurations. Only the trimer-descriptor GAP reproduces the distribution of the atomistic model over the whole range of bond angles.

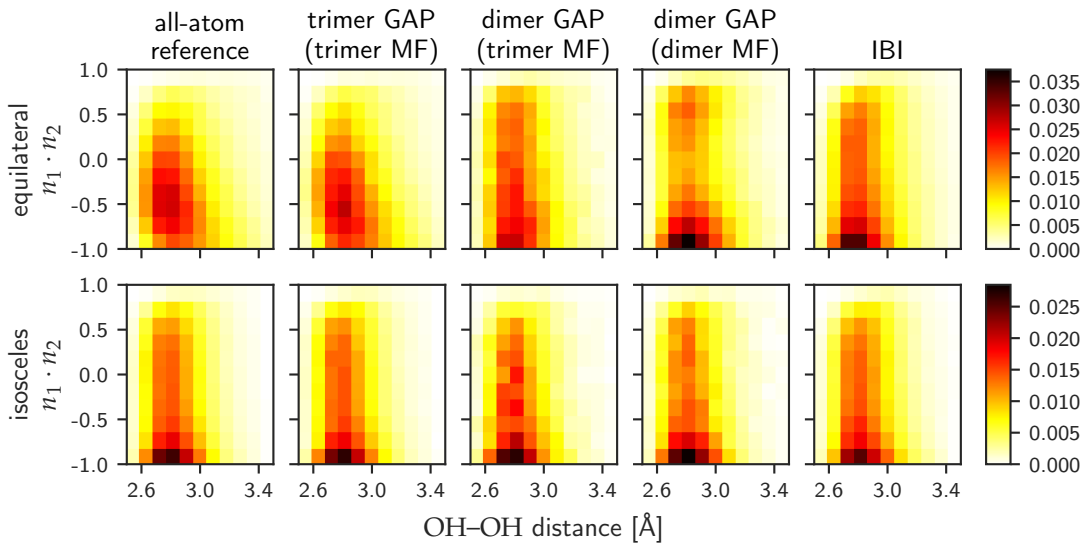

Figure 5.21: Methanol trimer: correlation of bond orientations. We measure the relative bond orientation between methanol molecules by the scalar product $n_1 \cdot n_2$ of the vectors which label the directions of the CG bonds. We show the distributions of relative bond orientation against OH–OH distance. These are plotted separately for equilateral configurations (top row) and isosceles configurations (bottom row). In the all-atom equilateral configuration, the OH sites point towards the centre of the triangle, hence they form angles of 120° with each other, corresponding to a dot product of -0.5. This is reproduced only by the trimer-descriptor GAP.

5.2.4. Tetramer

Atomistic structure

As in the cases of dimer and trimer, we employ half-open distance restraints between the centres of mass of the molecules to keep the cluster from dissociating. There are six restraints between all pairs of molecules, with the restraint setting in at 20 Å, and the force constant is 50 kcal/mol/Å^2. For IBI, we place the tetramer in a periodic cubic box with side length $L = 26$ Å.

In a system of four methanol molecules, configurations can contain tetramers, trimers, dimers, and free molecules. However, clusters of four methanol molecules predominantly form a square structure in which there are four hydrogen bonds linking all four molecules into a planar system. This can be seen in the ADF (Fig. 5.24 on page 121), where there are clear peaks at 90° and 2.7 Å (corresponding to the triplet formed by a corner of the square and its two hydrogen-bonded nearest neighbours) and at 45° and $\sqrt{2} \times 2.7$ Å ≈ 3.8 Å (corresponding to the triplet formed with one hydrogen-bonded nearest neighbour and the molecule diagonally across the square). The presence of other types of configurations is indicated by the smaller, secondary peak in the ADF at 60° and 2.7 Å that corresponds to configurations containing equilateral trimers.

CG potentials

GAP. As discussed in Section 4.2.2, a descriptor based on quadruplets of molecules would not be practicable. We already saw that the dimer descriptor is not sufficient to reproduce the methanol trimer distributions. Hence, we only consider GAPs that include the trimer descriptor.

We train a trimer-descriptor GAP from tetramer MF using the same hyperparameters as for the methanol trimer (Table A.2 in the Appendix); this includes the dimer core potential. The difference is only in the training data. Here, we use 12684 randomly sampled tetramer configurations. We note that it is possible to combine configurations from different systems in a single training set. Training a GAP from such a mixed set results in a potential that describes "average" molecular free energy contributions that are an approximation between the different monomer, dimer and trimer terms that best describe the PMF across systems. Here we include a GAP trained on a mixed set of trimer and tetramer MF that contained 12684 tetramer configurations (corresponding to 35508 triplets within the cutoff of the trimer descriptor) and 33380 trimer configurations. For comparison, we include the trimer GAP trained on trimer MF.

Pair potentials. As for the previous methanol trimer, we also present results for pair potentials obtained both by force-matching (using the same training data as for GAP) and by IBI (in a periodic box). The pair potentials are shown in Fig. A.2 in the Appendix.

Comparison of potentials

We show the force prediction performance of the different potentials, measured by the root mean squared error on a separate tetramer test set, in Table 5.3.

potential	derived from	RMSE / meV
monomer GAP	monomer MF	105
dimer GAP	dimer MF	53
trimer GAP	trimer MF	38
trimer GAP	trimer+tetramer MF	33
trimer GAP	tetramer MF	32
MSCG/FM	tetramer MF	59
IBI	tetramer RDF	67

Table 5.3: Methanol tetramer force prediction errors. Comparison of the force prediction root mean squared error (RMSE) on a separate mean force test set for different potentials.

MD results

As in the case of the methanol trimer, we separately show the distance distributions for vacuum/restrained simulations in Fig. 5.22 and for simulations using periodic boundary conditions in Fig. 5.23. The angular distribution functions shown in Fig. 5.24 clearly demonstrate that reproducing the pair distances does not, necessarily, imply a good reproduction of other measures of the structure.

We calculate the number of hydrogen bonds within the CG simulations based on an OH–OH cutoff of 3.4 Å. The fractional occupation of states with a given number of hydrogen bonds is shown in Table 5.4 and visualised in Fig. 5.25. Configurations with at least four hydrogen bonds describe a single bound tetramer. For these configurations, we show the distribution of angles and solid angles in Fig. 5.26.

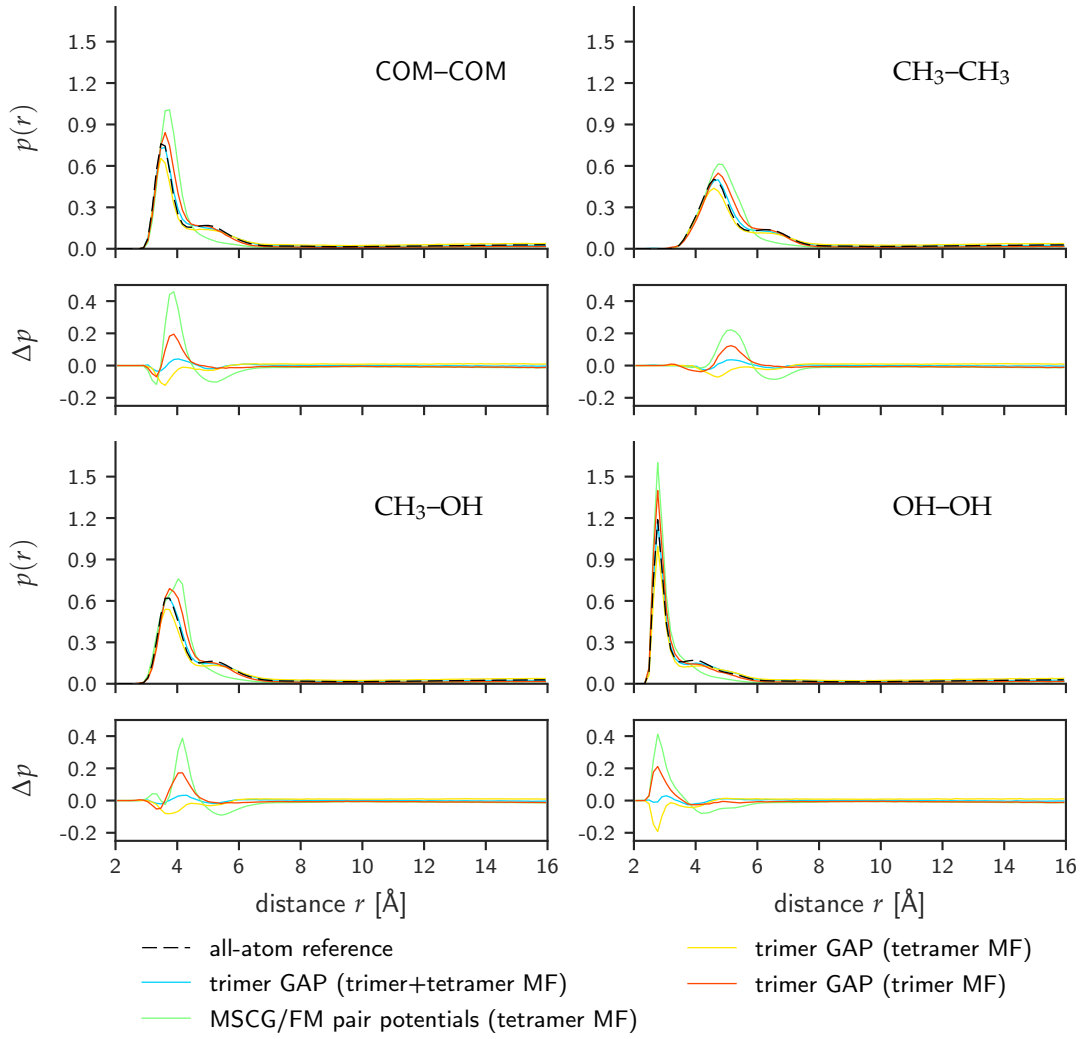

Figure 5.22: Methanol tetramer: distance distributions (for simulations carried out with distance restraints). COM and site–site distance distributions.

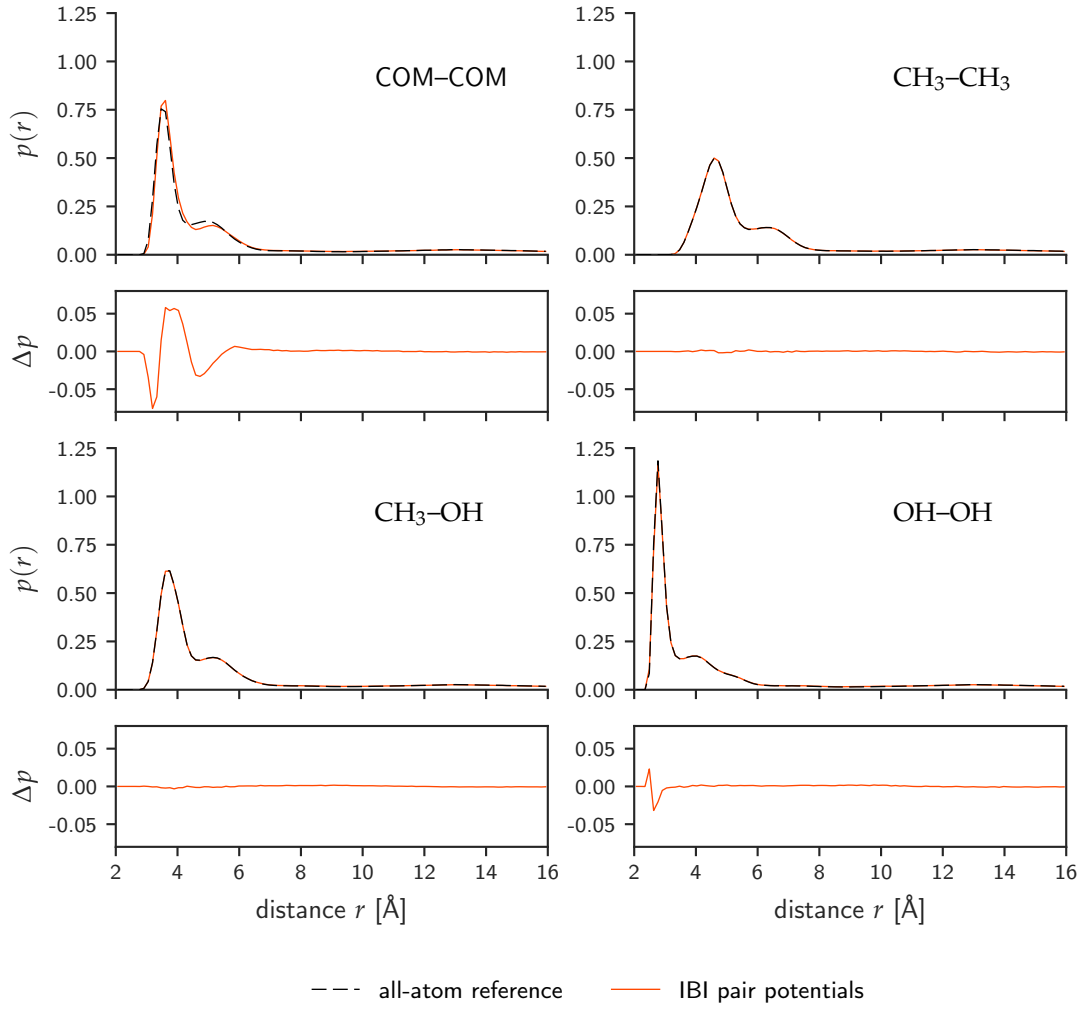

Figure 5.23: Methanol tetramer: distance distributions (for simulations carried out with periodic boundary conditions). COM and site–site distance distributions. By construction, the site–site distance distributions obtained for the IBI pair potentials match the all-atom distributions perfectly; the two lines are nearly indistinguishable. However, there is systematic discrepancy in the distribution of COM distances.

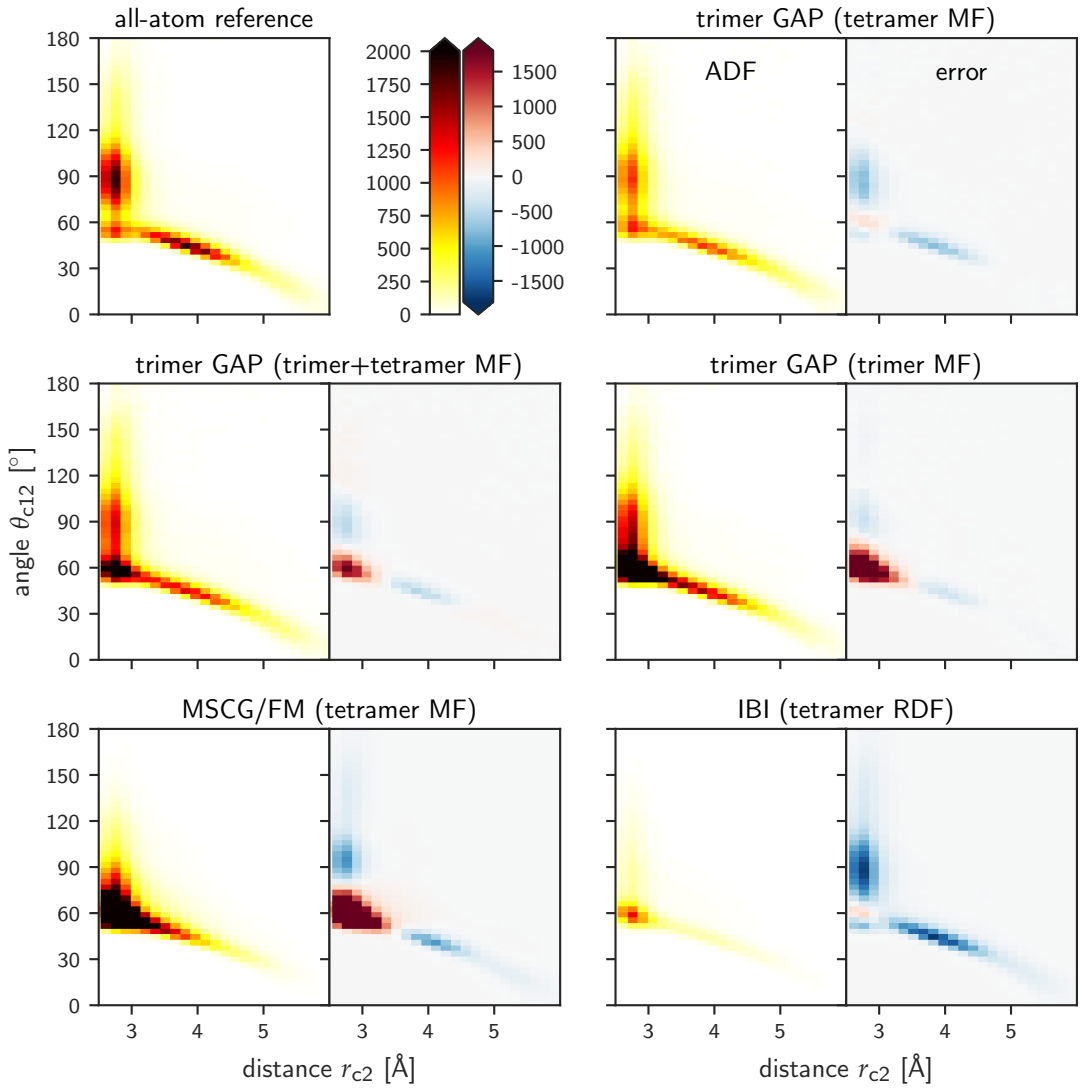

Figure 5.24: Methanol tetramer: Angular Distribution Function (ADF) for OH sites. The ADF (left-hand plot of each pair) is determined by considering all triplets of OH sites where the distance r_{c1} between a central site and its first neighbour is less than 3.4 Å. Here we show the 2D distribution of the angle θ_{c12} at the central site against the distance r_{c2} between the central site and its second neighbour. We only included configurations in the ADF calculation in which all pairwise COM distances were less than 10 Å. For the all-atom model, there is no discernible difference between the results for simulations using periodic boundary conditions and for simulations with distance restraints and hence we do not consider them separately in this figure. The right-hand plot of each pair shows the deviation from the all-atom result.

potential	derived from	number of hydrogen bonds / %						
		6	5	4	3	2	1	0
all-atom (restraints)		0.0	10.0	27.6	21.8	21.9	13.8	4.9
all-atom (periodic box)		0.0	11.1	30.8	23.6	18.8	11.1	4.5
trimer GAP	trimer MF	9.6	22.4	24.8	16.9	15.7	8.1	2.5
trimer GAP	trimer+tetramer MF	2.2	12.9	22.9	22.3	21.7	12.9	5.0
trimer GAP	tetramer MF	0.2	7.4	17.6	20.1	22.4	22.4	10.0
MSCG/FM	tetramer MF	23.9	29.0	16.3	12.5	6.6	6.8	4.8
IBI	tetramer RDF	4.4	18.5	20.4	18.6	13.5	12.9	11.7

Table 5.4: Methanol tetramer: relative frequencies of number of hydrogen bonds. We measure the number of hydrogen bonds in each configuration by counting the number of OH–OH distances that are less than 3.4 Å.

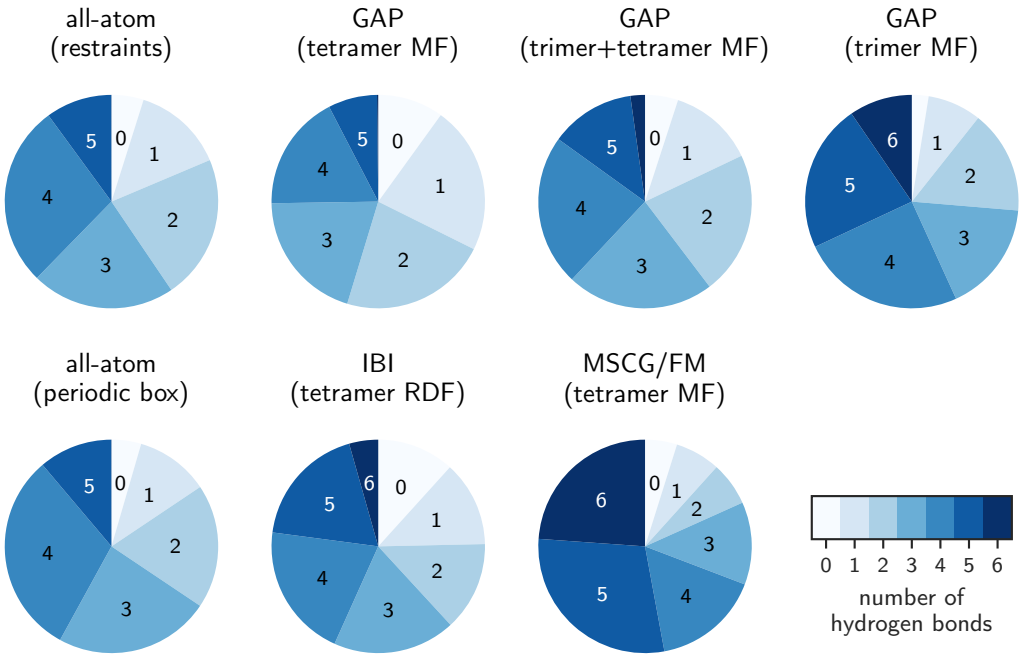

Figure 5.25: Methanol tetramer: relative frequencies of number of hydrogen bonds. We show the fractional occupation of states with different numbers of hydrogen bonds, measured by the number of OH–OH distances less than 3.4 Å. This is a visualisation of the values in Table 5.4.

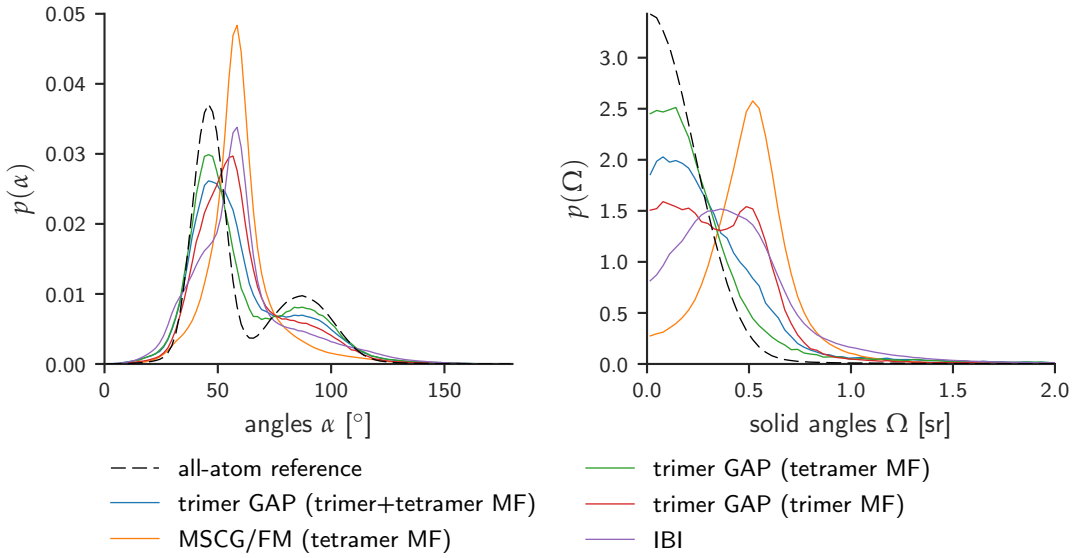

Figure 5.26: Methanol tetramer: angle distributions within bound configurations. Only configurations with at least four OH–OH distances less than 3.4 Å were included, corresponding to all four molecules linked together by hydrogen bonds. *Left:* distribution of angles between triplets of OH sites. *Right:* distribution of solid angles from each OH site to the others. A planar square structure corresponds to vanishing solid angles and two peaks at 45° and 90° degrees for the angles. A regular tetrahedron has solid angles of approximately 0.55 sr, and a peak at 60° for the angles. Whereas the trimer GAP reproduces the all-atom distribution to a large extent, both structure- and force-based pair potentials fail to describe the planar structure in the all-atom model correctly and tend to form tetrahedrons instead.

Unlike the cases of the methanol dimer and trimer, our GAP-CG simulations do not exactly reproduce the behaviour of the all-atom tetramer model. This implies that, in our two-site CG description, the PMF of the methanol tetramer cannot be partitioned exactly into monomer, dimer, and trimer terms. Hence, the GAP approach is no longer exact and, instead, becomes an approximation. This approximation might be improved by increasing the size of the training set.

However, even though the trimer-descriptor GAP does not reproduce the distance distributions as well as the IBI pair potentials, it can correctly reproduce the *structure* of four methanol molecules that form a planar square using four hydrogen bonds. The pair potentials completely fail to reproduce this structure, as is clearly demonstrated by the angular distribution functions. Overall, the trimer-descriptor GAP still provides a much more accurate representation of the distributions of the methanol tetramer as a whole than the other potentials.

This is true even for the trimer GAP trained on *trimer* MF. However, it puts too much weight on trimer configurations within the tetramer system and does not describe the square structure as well as the GAP trained on tetramer configurations. The reverse is true as well, that is, the GAP trained on tetramer MF does not describe the trimer system as well as the GAP trained on trimer MF. We note that, by combining training configurations from both tetramer and trimer systems, we can obtain a GAP that interpolates between both systems and finds the best separation into (effective) dimer and trimer interactions that gives reasonable results both for trimer and tetramer systems.

Generally, the best results will be achieved by training the GAP from data of the system in which it will be used. Here, we simply reused the vacuum dimer GAP as a core potential. In the general case, there will be effective contributions to the dimer/trimer interactions due to environmental influences (four- and higher-body effects such as contributions to the structure due to dense packing). Hence, for the bulk systems considered in the remainder of this thesis, we will not use a monomer or dimer core potential; instead, we will teach monomer, dimer, and/or trimer descriptors simultaneously.

5.3. Bulk Simulations

Whereas small clusters are dominated by direct interactions, in bulk systems environmental effects play a significant role. This was already discussed in relation to Boltzmann Inversion (see Sections 3.4.1 and 3.4.2), where an iterative procedure became necessary for an accurate reproduction of all-atom distributions due to correlations between different degrees of freedom. A large contribution to the overall structure simply comes from the dense packing of the molecules, determined by the balance between the external pressure and the avoidance of atomic overlap due to the strong core repulsion of atoms. This allows even a potential that only describes short-range repulsion of CG sites ("core-only potential") to already reproduce the all-atom structure to a large extent. The predominance of packing effects has several implications.

In structure-based approaches, the site–site pair potentials only need to provide a small adjustment to reproduce distance distributions, and this allows for a much better reproduction of orientational distributions in the bulk than in the clusters.

In force-based approaches such as GAP, the mean forces that we aim to reproduce are different from those in vacuum or clusters. Hence, we cannot simply reuse vacuum or cluster potentials.* The molecular contributions (dimer and trimer terms) now have to take into account the effective, environment-mediated interactions. This is demonstrated by the difference between vacuum and bulk mean forces along a cutline (shown in Fig. 5.29 on page 130). However, even though the molecular terms will be different between cluster and bulk, we can still consider a decomposition of the PMF into monomer, dimer, and trimer contributions.

To demonstrate the general applicability of the GAP approach to different systems, here we consider bulk models for methanol as well as for benzene. Whereas small clusters have very short decorrelation times and can be considered equilibrated almost instantly, in bulk systems careful equilibration becomes crucial. Moreover, to ensure there are no spurious effects due to interactions with the periodic mirror images of the system, we need to check that the properties of interest in the system are converged with respect to the box size. In Section 5.3.1 we outline our equilibration protocol and discuss convergence with respect to system size. We then show in Section 5.3.2 that, indeed, the mean forces only depend on a local environment, which is a necessary condition for a successful decomposition of the PMF into a cluster expansion of molecular terms with a

*In principle, we could use the cluster GAP as a core potential to describe the direct contribution to the effective forces in the bulk system, and teach the difference. However, this may not always improve the training, and in the following we only keep the repulsive short-range pair core potential for the non-bonded interactions between CG sites.

finite cutoff. We discuss the results of our CG simulations for methanol in Section 5.3.3 and for benzene in Section 5.3.4.

5.3.1. Simulation Protocol

In all atomistic simulations of the bulk systems, we use periodic boundary conditions, and employ a force field cutoff of 9 Å. The systems were equilibrated using the protocol given in Table 5.5. The trajectory root mean squared deviation from the initial structure, calculated using the minimum image convention, reached a plateau after 0.5 ns for methanol and after 1 ns for benzene, indicating that our equilibration times are sufficiently long.

	type	steps	(time)
energy minimisation	steepest descent	100	
	conjugate gradient	9900	
temperature ramp from 0 K to $T = 300$ K	NVT	50 000	(25 ps)
equilibration at $T = 300$ K	NVT	50 000	(25 ps)
pressure equilibration	NPT	3 000 000	(1.5 ns)
determine average volume	NPT	2 000 000	(1.0 ns)
equilibration at averaged volume	NVT	2 000 000	(1.0 ns)

Table 5.5: Equilibration protocol for bulk systems.

System size convergence

For bulk systems, it is important to check that our model is converged with respect to the system size.

For methanol, we considered initial configurations of $6\times6\times6 = 216$ and $7\times7\times7 = 343$ molecules. Following the equilibration procedure described in Table 5.5, this resulted in final box lengths of $L = 24.205$ Å and 28.226 Å, respectively. In each case, we carried out a 10 ns production run in the NVT ensemble. We compare RDFs from these final runs in Fig. 5.27(a) (see following page). The difference between the RDFs is negligible. Hence, in the following we only consider the system with 216 methanol molecules.

We did not carry out a separate system size test for benzene. However, in an earlier investigation we considered the convergence with respect to system size for imidazole, a planar molecule of similar size and structure to benzene. For the imidazole system, we considered initial configurations of $5\times5\times5$, $10\times10\times10$ and $15\times15\times15$ molecules.

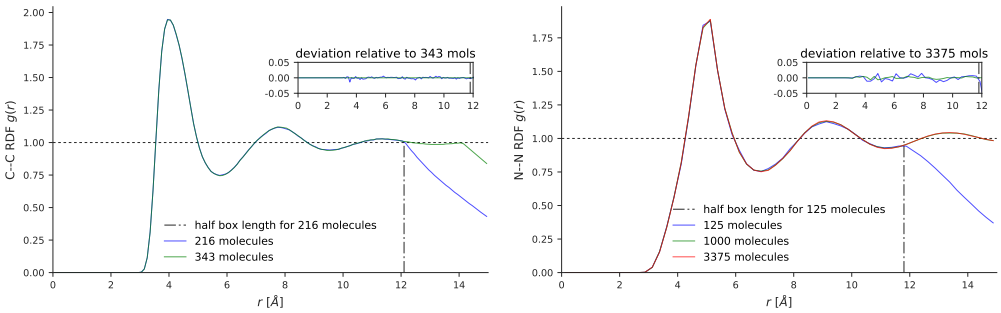

(a) Methanol **(b)** Imidazole

Figure 5.27: Convergence with respect to system size. *Left:* bulk methanol RDF between carbon atoms, for 216 and 343 molecules. *Right:* bulk imidazole RDF between nitrogen atoms, for 125, 1000, and 3375 molecules.

As can be seen in Fig. 5.27(b), even for $N = 125$ the results are sufficiently converged. As we used the same force field, we would expect very similar results for benzene. In the following, we use $N = 216$ for the benzene bulk simulations, corresponding to a starting configuration of $6 \times 6 \times 6$ molecules in a box with side length $L = 32.0205\,\text{Å}$.

5.3.2. Locality Test

We will show that the PMF is indeed a *local* function, in the sense that it can be written as a sum of local molecular contributions. We want to know to what extent it is possible to describe the (full) interactions of the system by only considering contributions associated with local environments. This is a critical requirement for our approach as any practicable CG interaction potential requires a cutoff.

To this end we consider a single molecule, referred to as the "central molecule"; its local or inner environment, given by all molecules within a cutoff radius from the central molecule (measured in terms of COM distances between the molecules), also called the "neighbourhood"; and the outer environment given by all other molecules.

For the locality test in methanol, we compare two cutoffs, 6.0 Å and 9.2 Å, corresponding to the first and second minima of the COM RDF, respectively. To take into account statistical uncertainty, we repeat the procedure discussed below for 10 different starting configurations. In these configurations, the smaller cutoff includes 12 to 17 molecules (average 14.0); the larger cutoff includes 44 to 55 molecules (average 50.5).

In the first stage, the central molecule and its inner environment are constrained on their CG sites, with all other molecules moving freely, and we run MD for 2 ns to sample different configurations of the outer environment. We take snapshots from these

constrained MD runs every 125 ps, corresponding to 16 different outer environments per inner environment. In the second stage, we constrain the CG sites of all molecules and run MD for 100 ps to calculate the mean forces on the CG sites due to the atomic fluctuations.

We now consider the mean forces on the CG sites of the central molecule for different environments. We denote these mean forces as $F_k^{\mathrm{MF}}(\gamma, \omega)$, where k indexes the six mean force components on the two sites of the central molecule, γ identifies the configuration of the inner environment, and ω identifies the configuration of the outer environment (for a given γ). For each inner environment, there is some variation in the mean forces due to the changes in the outer environment. This can be quantified by the deviation between the individual $F_k^{\mathrm{MF}}(\gamma, \omega)$ and the average value over the outer environments for a given inner environment. In Fig. 5.28 (on the following page) it can be seen that this mean force deviation is generally smaller for a larger cutoff.

Note that this locality test does not consider whether coarse-grained particles and forces are a good approximation to the atomistic interactions; it only investigates whether the CG forces due to a local neighbourhood can be a good approximation to the "converged" CG forces calculated in the full system. Thus we are addressing the question "How many neighbours do we have to include to appropriately describe the system?"

The spread we measure in a locality test for a given cutoff describes the maximum accuracy we can expect to obtain when separating the PMF into a function of local neighbourhoods based on this cutoff.

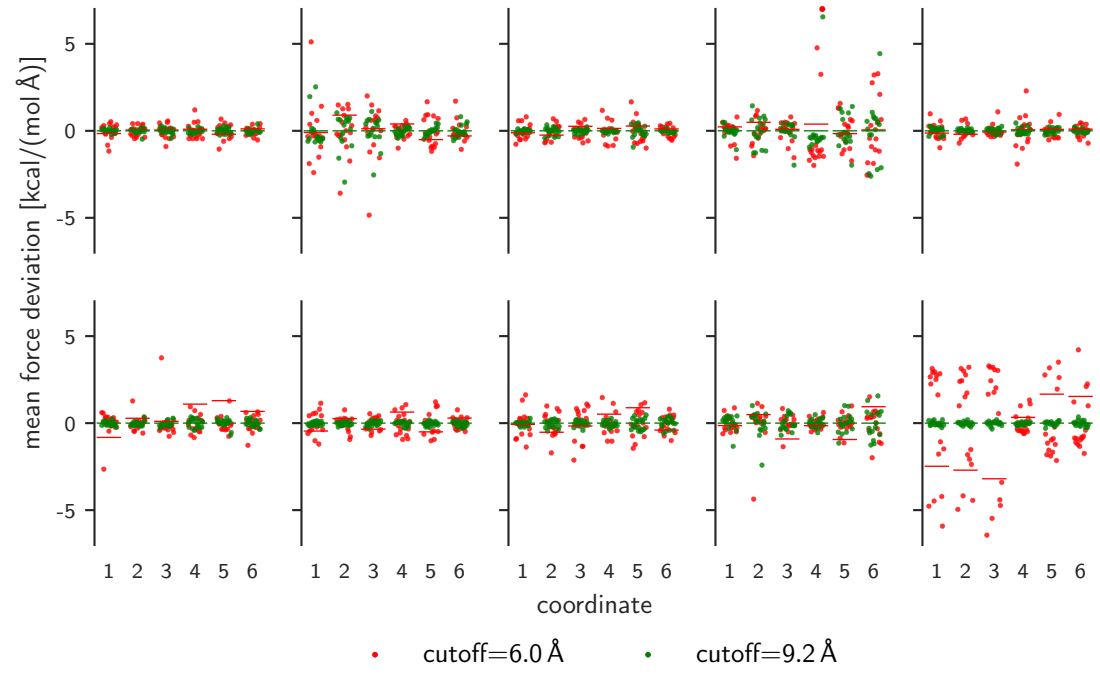

Figure 5.28: Locality test of the mean forces for methanol. Distribution of the mean force components (1 to 6) of the central molecule for different configurations beyond the neighbour shell, for two different sizes of the neighbour shell. The reference values are given by the average over outer environments for the larger 9.2 Å cutoff. Horizontal bars indicate the deviation of the *average* over outer environments for the smaller 6.0 Å cutoff. Each plot corresponds to one of the ten configurations of the inner environment.

5.3.3. Methanol

For the methanol monomer, the bond length distribution in bulk shows no significant difference from that of the vacuum monomer. Hence, we expect the monomer term for the bulk to be similar to the monomer GAP trained on vacuum grid MF. [Based on the difference in bond length average and the linear fit force constant, we would expect a deviation of about 0.015 eV/Å.]

While the monomer terms for methanol are very similar, intermolecular interactions in bulk are significantly different from the clusters due to the environmental contributions to the mean forces. This can be seen in the MF cutlines, which we compare for a dimer in bulk and for a dimer in vacuum in Fig. 5.29. Note that the bulk cutline goes to zero much faster than the vacuum cutline; this motivates the use of much smaller cutoffs for the bulk systems.

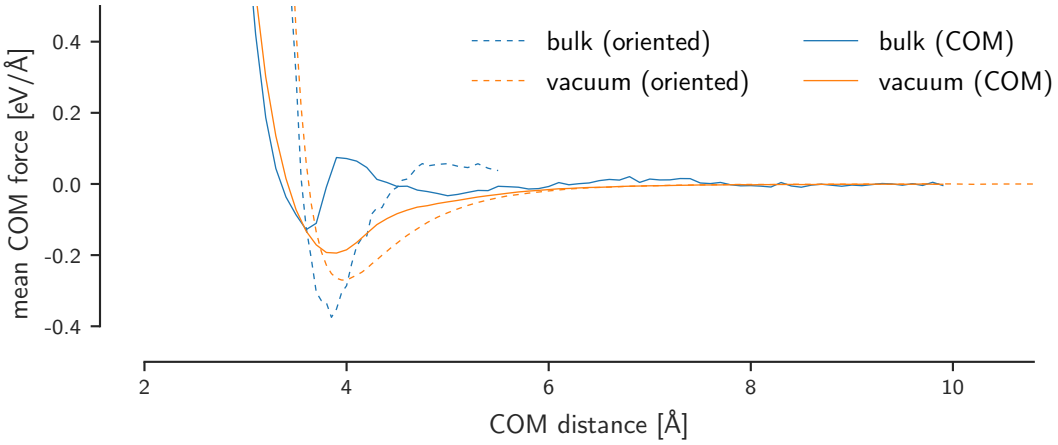

Figure 5.29: Comparison of COM force between bulk methanol and vacuum dimer. "Oriented" refers to a fixed relative orientation of two methanol molecules at a range of COM distances. "COM" refers to only the COM distance being fixed, while averaging over the orientations of the molecules.

CG Potentials

GAP. Because the monomer terms are very similar for bulk and for vacuum, we could use the vacuum monomer potential as a core potential for the bulk GAP potentials. When teaching just monomer and dimer terms simultaneously, there is indeed a slight systematic difference between the monomer term derived from bulk and the vacuum monomer. However, when teaching monomer, dimer and trimer terms simultaneously, the bulk-induced monomer term exactly corresponds to the vacuum monomer MF.

It would also be possible to use a GAP from the cluster expansion as a core potential; this would describe the *direct* contribution to the mean force, and we could teach a dimer and/or trimer descriptor on top of this core potential to capture the environmental contribution to the mean forces.

However, it may not always be feasible to carry out the cluster expansion, and here we show that we can learn monomer, dimer and trimer terms directly from force observations in the bulk system. We use a training set of 30 configurations of the CG coordinates. Due to the large number of collective variables, we only collected MF in runs of approximately 100 ps. As these are relatively short MD runs, the error in the mean forces is larger, and we set $\sigma_f = 5\,\mathrm{meV/Å}$ instead of $0.5\,\mathrm{meV/Å}$.

We consider two GAP potentials for bulk methanol. In the first one, we teach monomer and dimer descriptors simultaneously; we refer to this as "dimer GAP". In the second one, we teach monomer, dimer, and trimer descriptors simultaneously; we refer to this as "trimer GAP". In both cases we do not make use of cluster potentials as a core potential. However, we do keep the short-range repulsive pair potentials discussed in Fig. 5.5 on page 95. Compared to the methanol cluster GAPs, we reduce the dimer cutoff from 9 Å to 6 Å, and the trimer cutoff from 6 Å to 4 Å, corresponding to the second and first minima in the all-atom COM RDF, respectively. For all other hyperparameters we use the same values as for the methanol cluster GAPs (see Tables A.1 and A.2 in the Appendix).

Pair potentials. As in the case of the clusters, we consider both a force-matching and a structure-matching set of pair potentials. For the force-matching pair potentials, it turns out that there is no significant difference between a 6 Å and a 9 Å cutoff. Our force-matching pair potentials are equivalent to those derived by Izvekov and Voth [107], and we obtain the same COM and site–site RDFs in our simulations.

The structure-matching pair potentials are, again, obtained by Iterative Boltzmann Inversion using the VOTCA package, using the same settings as described in Fig. 5.16 on page 108.

Core-only potential. To demonstrate the significant contribution to the overall structure that is solely due to the short-range repulsion which is needed to prevent overlap of atoms, we considered two different "core-only" potentials. The MF core potential is the one used to provide the short-ranged repulsion for non-bonded interactions in the methanol clusters (*cf.* Fig. 5.5). The DBI core potential is based on a spline fit to the Direct Boltzmann Inversion of the short-range part of the all-atom RDFs. Beyond their

respective short-range cutoffs, the core potentials are exactly zero.

The pair potentials are shown in Fig. A.3 in the Appendix. We compare the COM force predictions along a cutline in Fig. 5.30. In Table 5.6 we compare the RMSE of the predictions of the various potentials on a test set of 10 configurations.

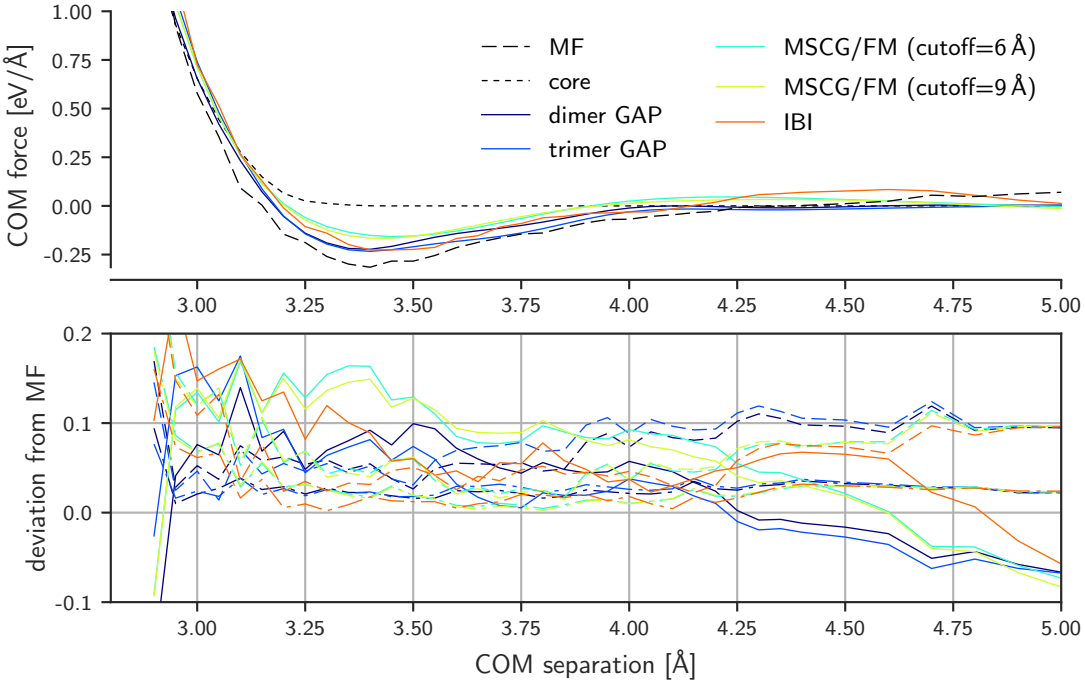

Figure 5.30: Bulk methanol: forces along a dimer cutline. Comparison of CG force predictions with bulk MF. For a more detailed explanation, see Fig. 5.7 on page 97.

potential	derived from	RMSE / meV/Å
monomer GAP	monomer MF	157
dimer GAP	dimer MF	84
dimer GAP	trimer MF	82
trimer GAP	trimer MF	170
IBI pair potential	bulk RDF	96
MSCG/FM pair potential	bulk MF	90
dimer GAP	bulk MF	70
trimer GAP	bulk MF	68

Table 5.6: Bulk methanol force prediction errors. Comparison of the force prediction root mean squared error (RMSE) with respect to MF in a separate test set of 10 bulk methanol configurations for different potentials.

MD results

We show the COM and site–site RDFs for simulations using the different potentials in Figs. 5.31 to 5.34. We show the angular distribution function of the centres of mass of the molecules in Fig. 5.35. In Figs. 5.36 and 5.37 we show distributions of the relative orientation between pairs of methanol molecules as a function of their distance. We consider both the dihedral angle formed by the two CG bonds as well as the dot product of the vectors indicating their direction.

The core-only potentials already reproduce the all-atom RDFs to a large extent, with the exception of the sharp first peak in the OH–OH and COM RDFs that is due to the hydrogen bonds between OH sites. Clearly, the force-matching pair potentials do not reproduce the structure as well as the IBI pair potentials. Here, the dimer GAP shows some overstructuring compared to the IBI pair potentials; this is especially prominent in the peak at 60° and 3 Å in the ADF. However, when including molecular trimer interactions in the GAP, it can reproduce the all-atom orientational and angular distributions better than the IBI pair potential, despite being trained on forces.

Overall, the results in this section demonstrate that we can capture the full many-body Potential of Mean Force of the bulk methanol system to a high degree of accuracy using GAP, the resulting potential being able to reproduce the structure even better than structure-matching pair potentials.

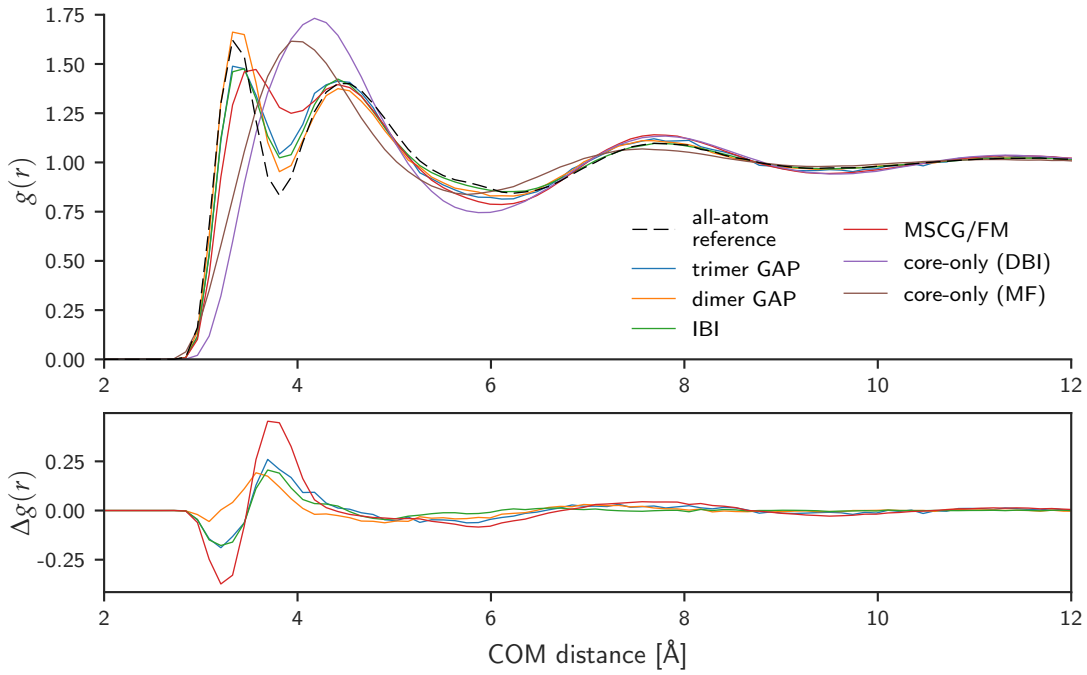

Figure 5.31: Bulk methanol: COM RDF for different potentials.

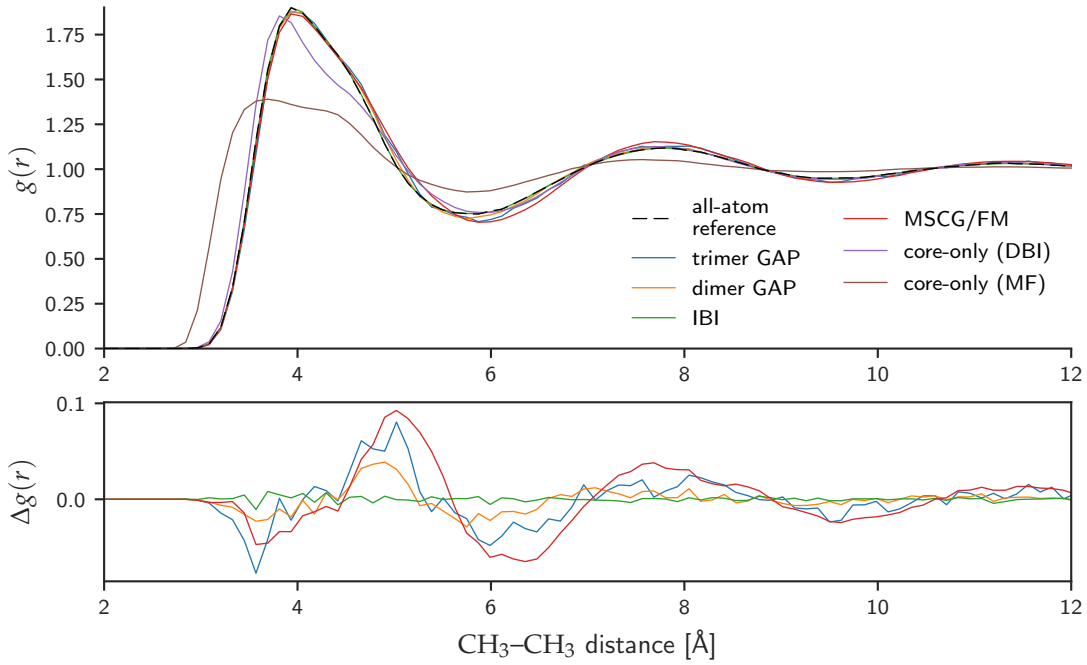

Figure 5.32: Bulk methanol: CH_3–CH_3 RDF for different potentials.

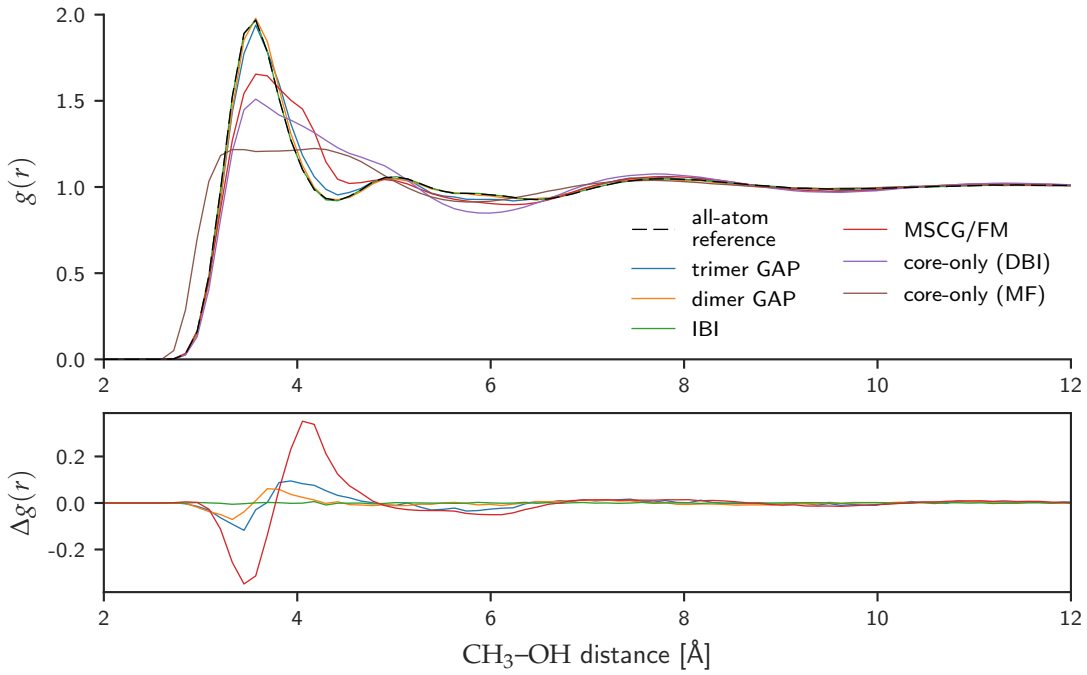

Figure 5.33: Bulk methanol: CH_3–OH RDF for different potentials.

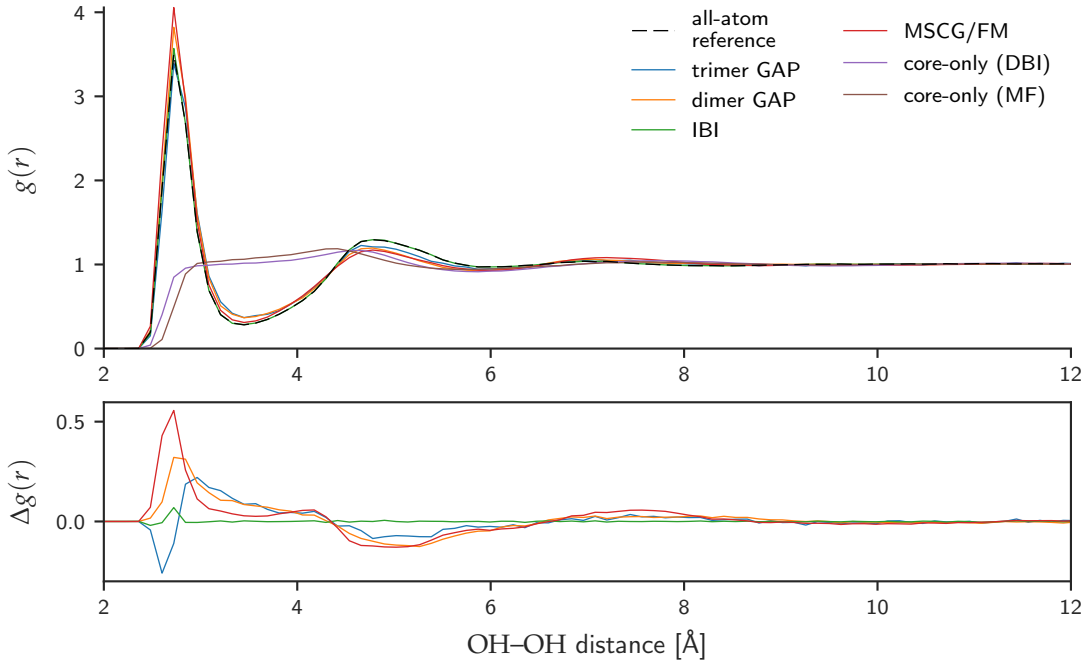

Figure 5.34: Bulk methanol: OH–OH RDF for different potentials.

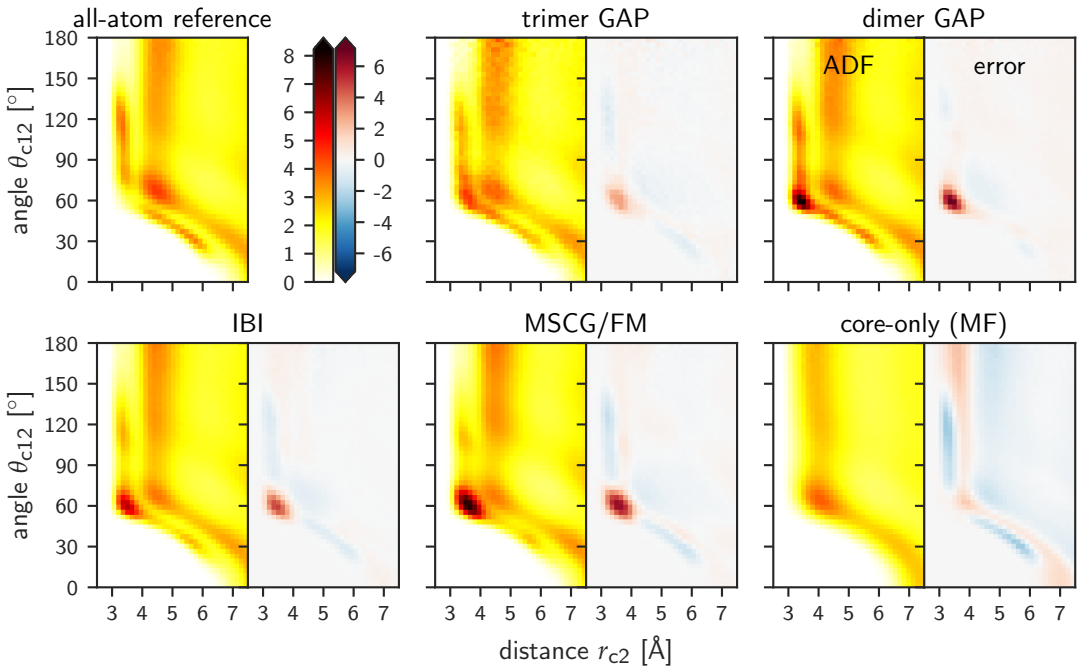

Figure 5.35: Bulk methanol: Angular Distribution Function of molecule COMs for different potentials. For a more detailed explanation, see Fig. 5.24 on page 121.

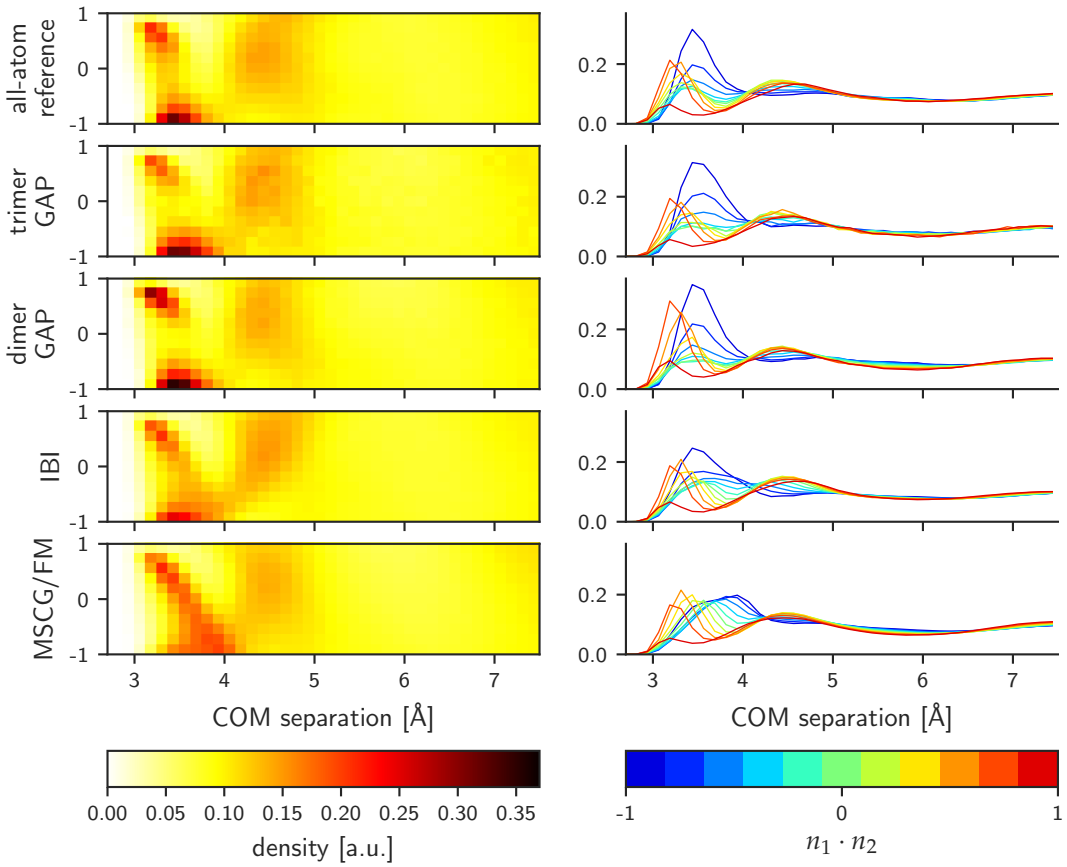

Figure 5.36: Bulk methanol: orientational distributions (1/2). This shows the distribution of the dot product between the normalised vectors that correspond to the directions of the CG bonds as a function of the COM separation between pairs of methanol molecules. *Left:* 2D distribution. *Right:* horizontal slices through this distribution. The distributions have been normalised along the distance axis according to the RDF convention.

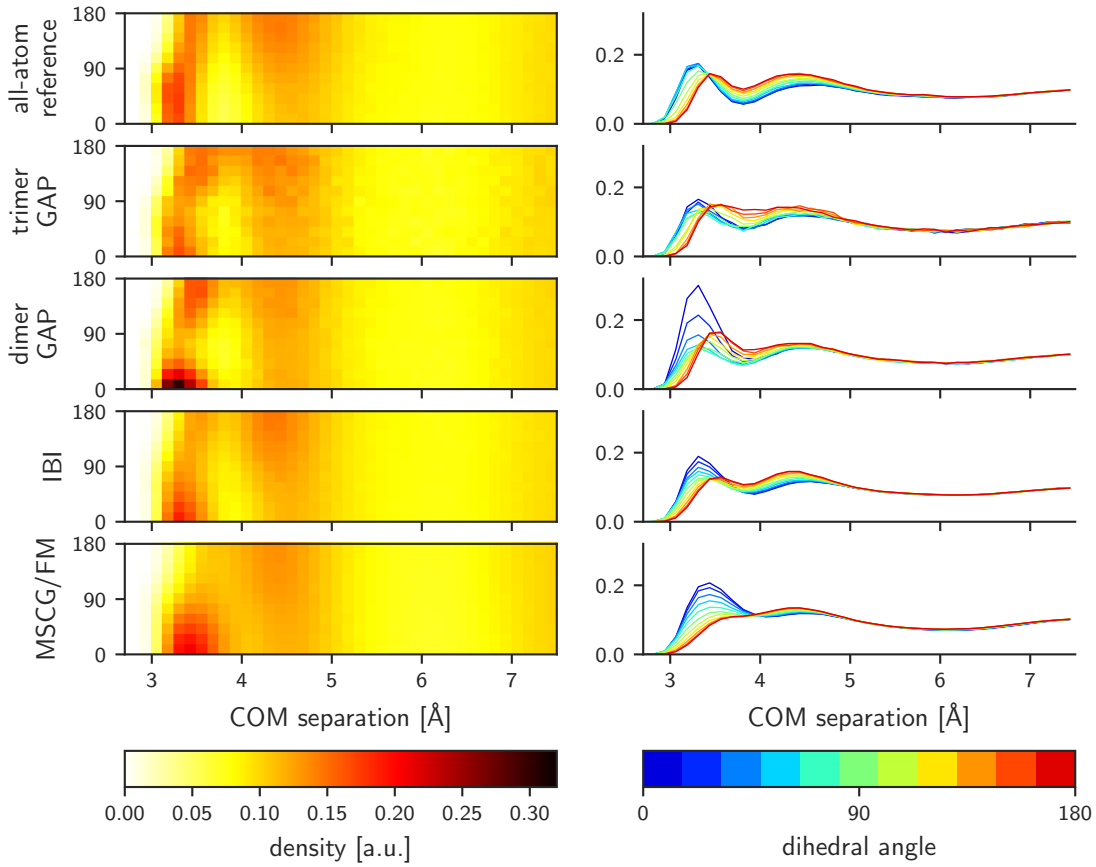

Figure 5.37: Bulk methanol: orientational distributions (2/2). This shows the distribution of the dihedral angle between the CG bonds as a function of the COM separation between pairs of methanol molecules. *Left:* 2D distribution. *Right:* horizontal slices through this distribution. The distributions have been normalised along the distance axis according to the RDF convention.

5.3.4. Benzene

We now consider bulk benzene, to demonstrate that the GAP approach can be easily applied to different systems. We described the molecular structure of benzene in Section 5.1.2 on page 86. In Table 5.5 we discussed the convergence with respect to the size of the periodic box for imidazole, a molecule very similar to benzene, and hence chose to simulate a system with 216 benzene molecules in a periodic box. Following the equilibration protocol (Table 5.5), we obtain a final box length of $L = 32.0205\,\text{Å}$.

CG Potentials

For benzene, we derive a core potential by fitting pair potentials to the bulk mean forces using a short-range cutoff of 4 Å. The resulting potential is shown in Fig. A.4 in the Appendix. Note that the exact form of the core potential does not matter as long as it prevents overlap of CG sites and does not extend significantly into regions that are considerably populated.

GAP. In the case of benzene we do not go beyond dimer interactions, as we will see that this already describes the full interactions sufficiently well. Moreover, the trimer descriptor would be too computationally expensive for bulk benzene. We teach both monomer and dimer simultaneously. The hyperparameters used in this case are similar to those used for methanol; the main difference is the covariance cutoff, which is now 8 Å, corresponding to the first minimum of the benzene COM RDF. The hyperparameters for the GAP used in MD simulations are given in Table A.3 in the Appendix. We will discuss the influence of different hyperparameters and how changing them affects the force predictions in Section 5.4.

Pair potentials. We generate pair potentials using both the force-based MSCG/FM approach and the structure-based IBI approach, as we did for methanol. The resulting pair potentials are shown together with the core-only potential in Fig. A.4 in the Appendix.

MD results

We show the COM and site–site RDFs in Fig. 5.38 and the COM ADF in Fig. 5.41. We consider the orientational distributions of pairs of molecules in Figs. 5.39 and 5.40.

The MD results demonstrate that the GAP can accurately capture the oscillations in the benzene pair orientations, even beyond the interaction cutoff of 8 Å. In contrast, the pair potentials only manage to describe the orientation up to about 6 Å, even though their cutoff is 10 Å.

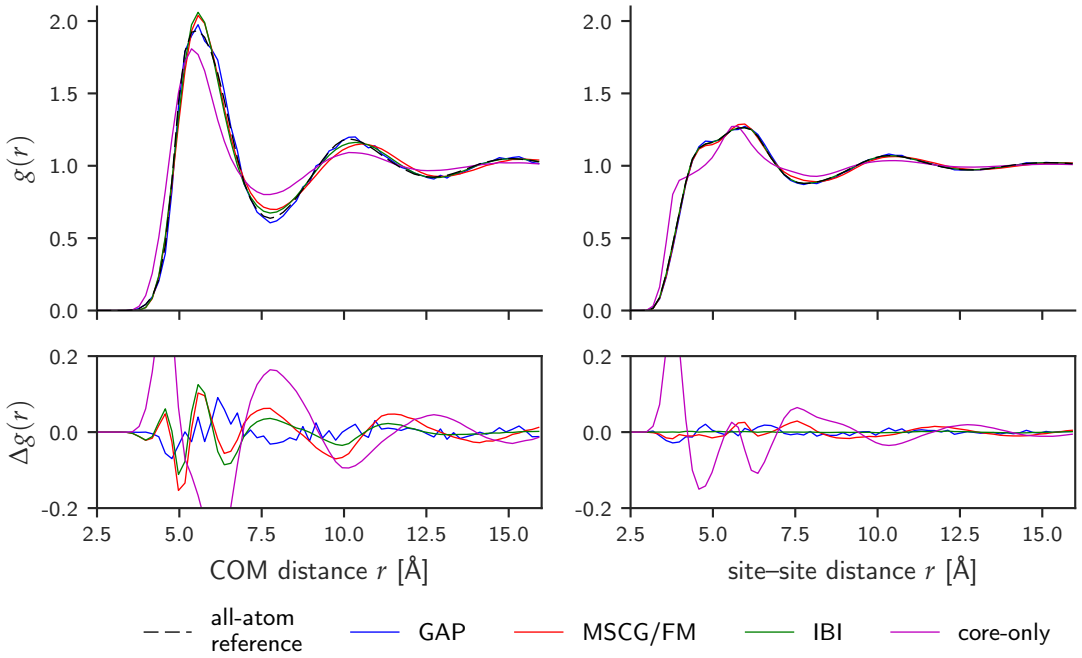

Figure 5.38: Bulk benzene: Radial Distribution Functions for COM and CG sites.

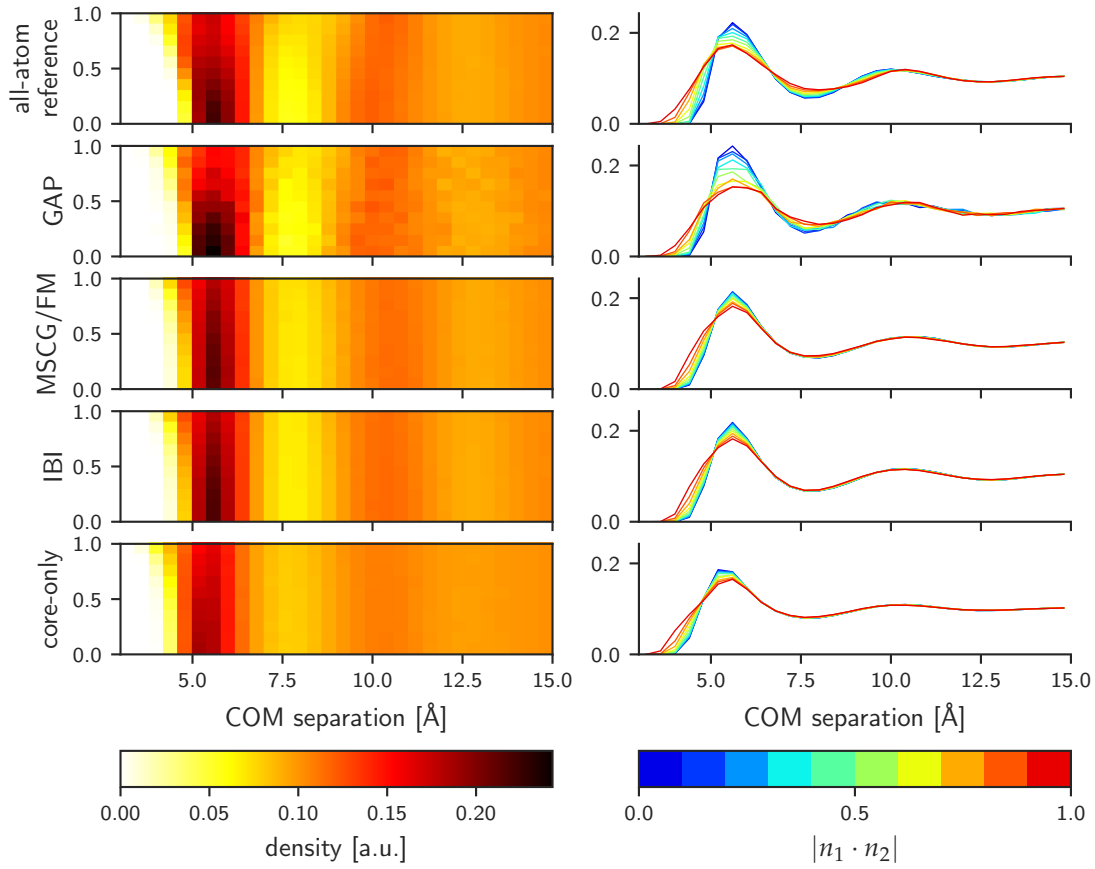

Figure 5.39: Bulk benzene: orientational distributions (1/2). This shows the distribution of plane orientations for pairs of benzene molecules as a function of their COM separation. The relative plane orientation is measured by the dot product $|n_1 \cdot n_2|$ between plane normal vectors; due to the symmetry of benzene, only the absolute value matters. A value of 1 corresponds to parallel orientation; a value of 0 corresponds to perpendicular orientation. *Left:* 2D distribution. *Right:* horizontal slices through this distribution.
This is an average over all pairs of molecules in the simulation box. The distributions have been normalised according to the RDF convention along the distance axis. Uniformly random orientations would lead to a uniform profile of the dot product, corresponding to all lines in the right-hand plot lying on top of each other.
Note that the GAP model describes the oscillations in the distribution as well as the all-atom model over the whole range of COM separations. This can still be seen at a COM distance of 10 Å, even though the (COM-based) GAP cutoff used in this model is only 8 Å. In contrast, the IBI pair potential that reproduces the site–site RDF perfectly has a cutoff of 10 Å, but it does not describe the oscillations in the orientational distribution beyond the first peak at 6 Å. This demonstrates that GAP based on the many-body dimer descriptor can reproduce the all-atom distribution much better than the pairwise spline-based potential.

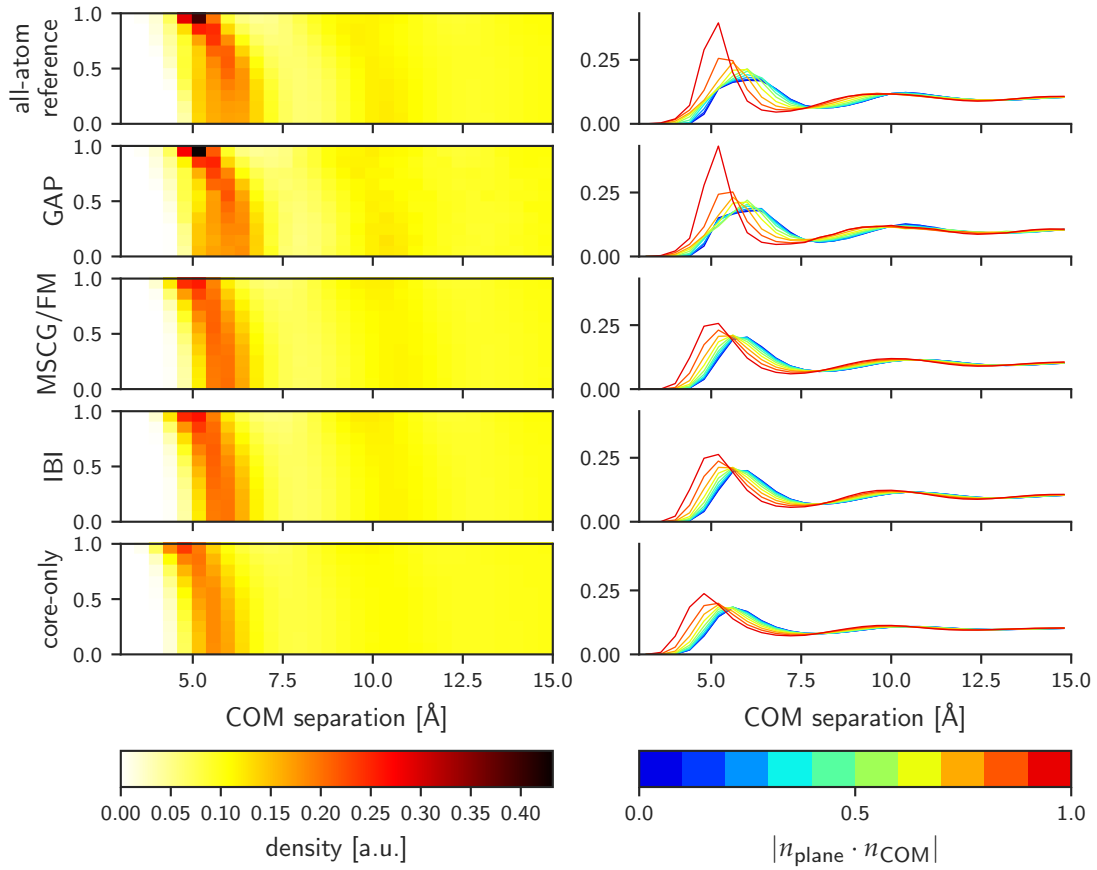

Figure 5.40: Bulk benzene: orientational distributions (2/2). Here we consider the distribution of the absolute value of the dot product between the plane normal vector n_{plane} of one molecule and the normalised vector along the COM separation axis n_{COM} to another molecule, as a function of their COM separation. This figure is otherwise analogous to Fig. 5.39.

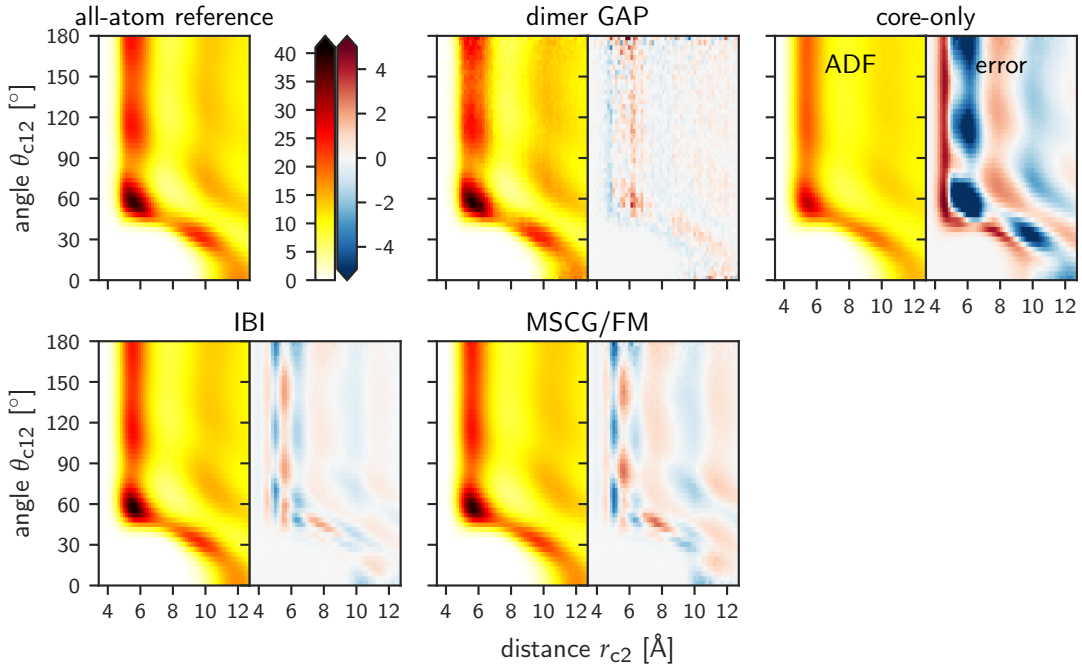

Figure 5.41: Bulk benzene: Angular Distribution Function (ADF) of the centres of mass of the benzene molecules, for the all-atom model and the different CG potentials. The ADF (left-hand plot of each pair) is determined by considering all benzene triplets in which the COM distance r_{c1} between a central molecule and its first neighbour is less than 7.5 Å. Here we show the 2D distribution of the angle θ_{c12} between the central molecule and its two neighbours against the distance r_{c2} between the central molecule and its second neighbour.
Even though the error is large, the core-only potential already reproduces the all-atom ADF qualitatively. Both IBI and MSCG/FM pair potentials show some systematic, patterned deviation from the all-atom ADF. The results for the dimer-descriptor GAP contain more noise, as the GAP-CG MD in this case was very expensive and we only obtained limited sampling. However, the GAP results show much less *systematic* deviation than the pair potentials.

GAP can recover higher symmetry than descriptor

Another demonstration of the power of the GAP approach and the flexibility of the potentials it can learn is the recovery of higher symmetry than that prescribed by the descriptor.

Benzene has six-fold symmetry; rotating a molecule by an integer multiple of 60° around the plane normal axis brings it into a state completely equivalent to the previous one. However, our three-site CG description of the benzene molecule only has three-fold symmetry: one would have to rotate a molecule by an integer multiple of 120° to return it to an equivalent state. However, despite this limitation, the GAP successfully learns that the three-site description is of a molecule with six-fold symmetry. This can be seen in the oscillation of the potential landscape of the dimer descriptor as a function of the in-plane rotation angles of both molecules shown in Fig. 5.42.

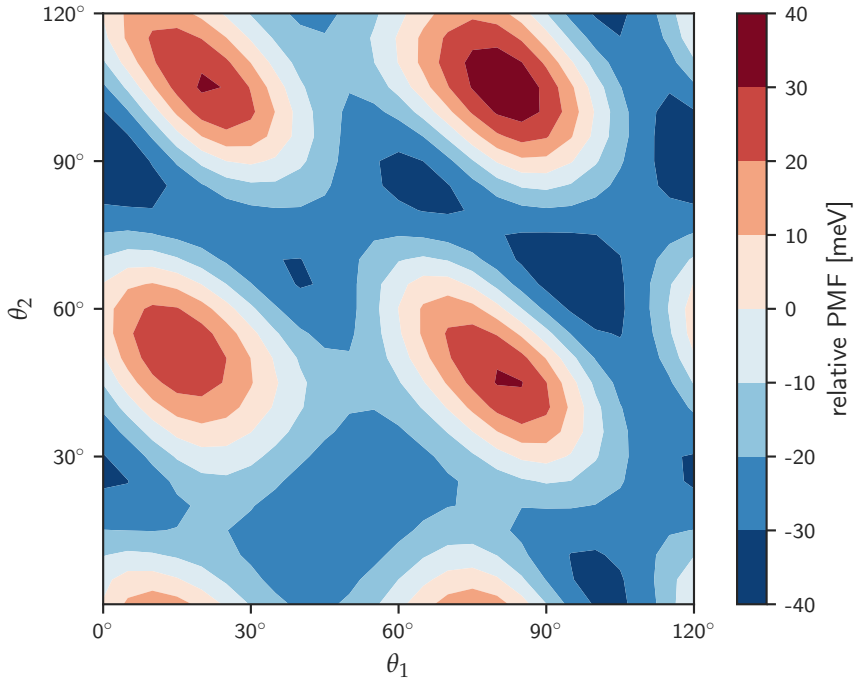

Figure 5.42: GAP-predicted PMF as function of rotation around plane normal axes for a dimer descriptor trained on bulk benzene MF. This reproduces the six-fold rotational symmetry of the all-atom description, even though the CG description only prescribes three-fold rotational symmetry: rotating either molecule by 30° around its plane normal axis leads to an (approximately) equivalent point in the GAP-predicted free energy surface.

5.4. Influence of Hyperparameters, Sparsification, and Training Data

So far, we have simply stated the values for hyperparameters and other settings used in generating GAP-CG potentials. Here, we discuss the influence of the hyperparameters (Section 5.4.1), the number and selection of sparse points (Section 5.4.2), and the size and noisiness of the training set. This is applied to the bulk benzene model presented in the previous section.

5.4.1. Hyperparameters

The squared exponential kernel discussed in Section 4.1.3 has several hyperparameters that influence the behaviour of the Gaussian process and hence the prediction performance (*cf.* page 52). The covariance cutoff, length scale and energy scale are properties of the *system* we want to model and are independent of the training data and its noise.

The choice of the covariance cutoff can be based on minima of the RDF; this minimises the number of interactions at the edge of the cutoff, which might be affected by cutoff effects. Whether only the first or also the second neighbour shell (or beyond) needs to be included is directly related to the locality of the mean forces, and we discussed this for bulk methanol in Section 5.3.2.

The length scale θ of the the tail of non-bonded interactions at longer range is on the order of 1 Å to 2 Å; for the more rapidly changing part of the interaction at short range or for stiffer bonded interactions, a length scale of around 0.3 Å is more appropriate. This variation in length scale necessitates the use of a "transfer function" to be able to describe both regimes at once (*cf.* Section 4.3.2). We demonstrated this for the methanol dimer in Fig. 5.5. In Fig. A.6 in the Appendix we show the effect of varying the length scales and the parameters of the transfer function.

The hyperparameter σ_f for the noise level in the force observations depends on the training data. When training from mean forces, we can calculate the error within the mean force calculations. In our case, the mean force is the mean of all constraint force samples, and we can measure the error based on the standard deviations of the constraint forces. We will discuss the much larger noise that is present when teaching from instantaneous collective forces in Section 5.4.3.

It should be noted that another source of noise from the point of view of the Gaussian process regression is a mismatch between the true function that we want to infer — in this case the PMF as a function of all $3N$ CG coordinates — and what is representable

by the descriptors used in the GAP model. If the descriptors cannot represent the true shape of the PMF, the noise level hyperparameter would need to be increased. Fortunately, our work suggests that the molecular descriptors based on monomers, dimers, and trimers already represent the PMF to a high degree.

The energy scale δ is hard to determine when we are only teaching from forces. In the simplest case, it is only the ratio δ/σ_f that matters. In Fig. 5.43 we show the force prediction error (relative to reference mean forces for a separate test set of 10 CG configurations that were not part of the training set) as a function of the energy scales for monomer and dimer descriptor, δ_{mono} and δ_{dim}, for different values of σ_f.

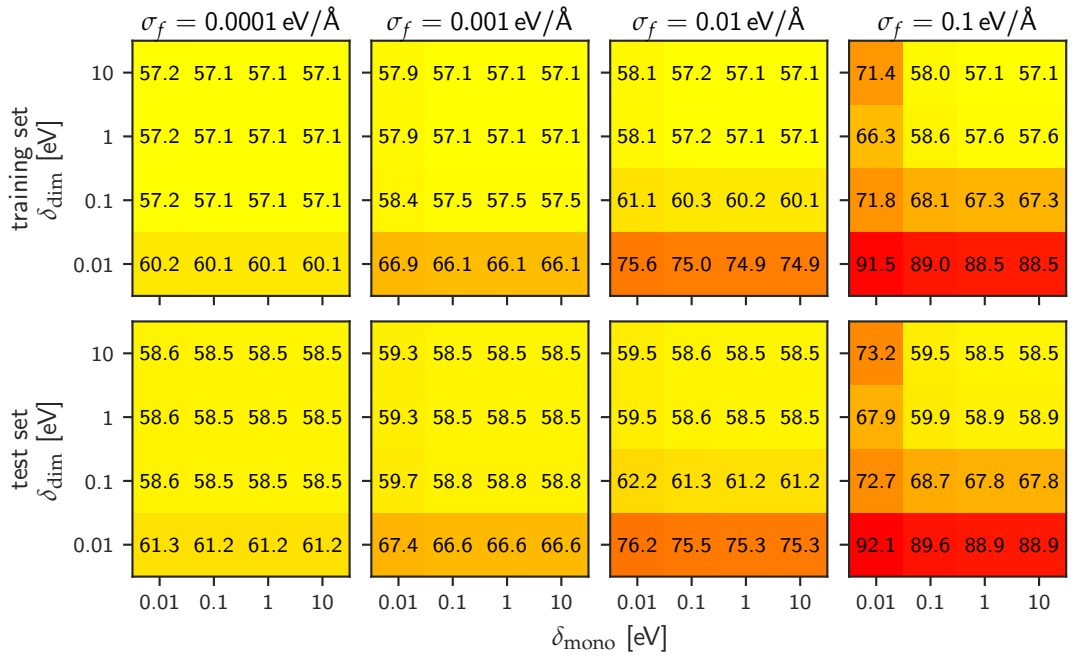

Figure 5.43: Influence of hyperparameters on GAP prediction: energy and noise scales. Force RMSE in meV in test set and training set for different combinations of the monomer and dimer energy scales (δ_{mono} and δ_{dim}, respectively), for different values of the noise level hyperparameter σ_f. All other settings as given in Table A.3. A large discrepancy between test set error and training set error indicates overfitting.

Finally, an important point to appreciate is that the more training data is available, the less is the influence of the prior as determined by the covariance function and its hyperparameters. Hence, with a sufficient number of training configurations, GAP predictions are robust with respect to changes in the hyperparameters.

5.4.2. Sparsification

Sparsification is a further approximation in the GAP approach to reduce the memory and computing requirements. The number of sparse points should be high enough to be able to represent the interaction potential accurately. In the limit of the number of sparse points being equal to the number of training points, there is no distinction between different methods for selecting the sparse points. However, in practice, the number of sparse points needs to be significantly smaller than the number of training points. In Fig. 5.44 we compare the force prediction error for the different selection algorithms introduced in Section 4.2.5 as a function of the number of sparse points.

The sparse points should evenly cover that part of the descriptor space that is actually explored by the system. A uniform selection of sparse points is the most favourable choice for low-dimensional systems and when a sufficiently large number of sparse points can be used. It avoids selecting sparse points that are close to each other, which would give rise to a high condition number of the covariance matrix. Interestingly, uniform selection still appears to perform best for the relatively high-dimensional benzene dimer descriptor (15 dimensions). For more drastic sparsification (smaller number of sparse points or higher descriptor dimension), we would expect CUR decomposition to outperform uniform selection.

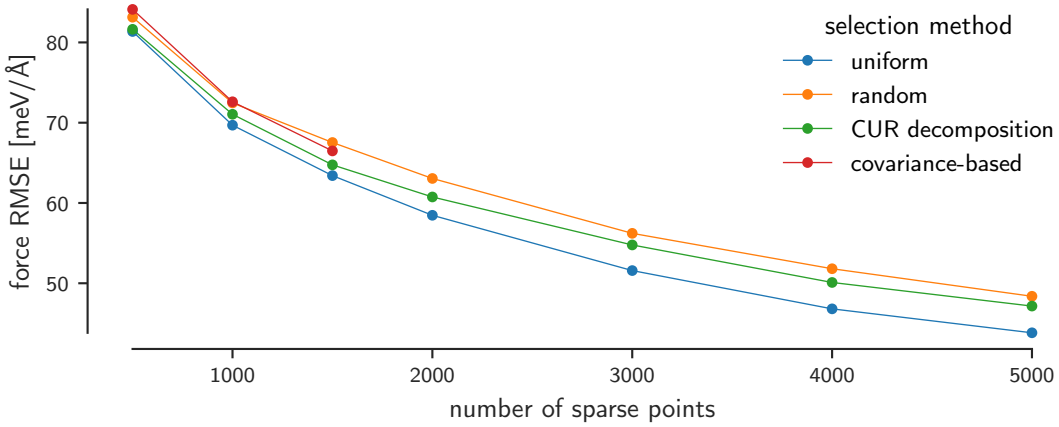

Figure 5.44: Influence of number and selection method of sparse points on GAP prediction. Force RMSE as a function of the number of sparse points, for different selection methods. All other settings as given in Table A.3. Covariance-based sparse point selection was too slow to be feasible for higher numbers of sparse points. Note that adding sparse points never reduces prediction performance, but yields diminishing returns.

5.4.3. Noisy Training Data

Throughout the present work, we train GAPs from mean forces (MF), as in this case the training data has very little noise (if sufficient sampling has been used). This allows us to focus on the task of reproducing the high-dimensional PMF from gradient observations. This approach amounts to a separation between *averaging* instantaneous forces to obtain the mean forces and *integrating* the mean forces to obtain the PMF. However, in practical applications it can be difficult to obtain mean forces with small errors due to the cost of calculations. Recovering a function from *noisy* observations is obviously more complicated than from data with no or only little noise. However, we showed that GAP can also be trained on samples of the instantaneous collective force (ICF), rather than the mean force, and thus carry out the averaging within the Gaussian process regression. The results shown in Fig. 5.45 (see following page) demonstrate that with increasing volume of training data the ICF-trained GAP predictions converge towards those of the MF-trained GAP.

In Section 5.2.1 we derived the appropriate noise level for teaching from ICF from the training data itself. In principle, this is also possible in higher dimensions. Note that we cannot simply take the *global* standard deviation of force components. Instead, we have to transform the force observations — that is, observations of gradients of the PMF with respect to positions of CG sites — into gradients of the local free energy contributions in the space of the descriptors that we want to train. We then need to construct a grid in the descriptor space and calculate the standard deviation of the gradients *in each bin*. The resulting value for the anticipated noise of the gradients in descriptor space would then have to be transformed back into Cartesian space. This is, however, beyond the scope of this thesis.

Importantly, it has been shown in the case of few dimensions that a reconstruction of the free energy surface from ICF using Gaussian processes is a more efficient use of computational resources (in terms of time-to-solution) than collecting less noisy mean forces for a smaller amount of configurations (*e.g.*, on a grid) [65, 66]. In the present work, we show that this is also true in the high-dimensional case of generating a coarse-graining GAP. In Table 5.7 we compare the predictive power of averaging within the GAP and external to the GAP based on a constant "total information content" of the training set. We consider ICF, constraint force samples, and the average over constraint force samples (*i.e.*, the mean force). In each case, the training set is based on 2000 samples. For ICF, these are samples for 2000 different configurations. Of these CG configurations, 20 were used to constrain the CG sites in Constrained MD runs

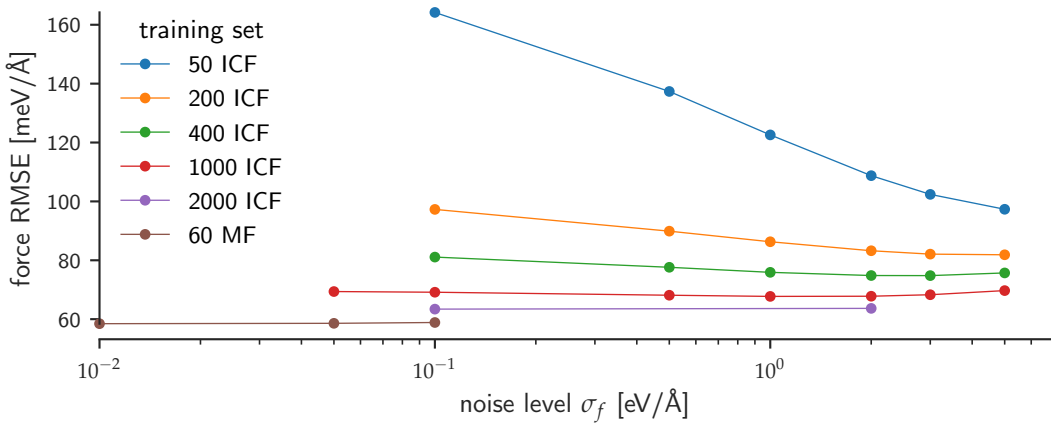

Figure 5.45: Influence of number of training configurations and noise scale on GAP force prediction in the case of noisy training data (ICF). We show the force prediction RMSE of GAPs taught from ICF, for different sizes of the training set and for a range of values of the noise level hyperparameter σ_f. All other settings as given in Table A.3. For comparison, we include the results for GAP trained on converged MF as before (60 configurations). This demonstrates that, with sufficient training data, teaching from ICF converges to the teaching from MF.
Importantly, this figure suggests choosing a very high noise level for a small and noisy training set. With an increasing number of training points, the RMSE curve flattens out. This means that the precise value of the noise level matters less. Note that a flat curve implies that the minimum is not stable and, hence, it does not make sense to "optimise" the noise level parameter; this would risk overfitting of the training data. As the prediction performance is stable with respect to the noise level for larger training sets, we can simply choose a sufficiently high noise level.

of 100 ps, during which we collected 100 constraint force samples per configuration. This results in a training set of 2000 frames. We separately calculate the average over these 100 samples, resulting in (comparatively noisy) mean forces for 20 configurations. For comparison, we also include the results for a GAP trained on the converged mean forces obtained at the end of the Constrained MD runs (corresponding to an average over 2×10^5 constraint force samples). The force prediction error depends on the value of the noise level hyperparameter. Teaching from an average of constraint force samples with a smaller noise level hyperparameter is equivalent to teaching from the samples directly with a higher noise level hyperparameter.* Here, training the GAP from ICF performs significantly better than training from an equivalent number of constraint force samples. Mean forces need to be converged to a much higher accuracy to provide better results.

	σ_f / eV/Å					
	0.001	0.01	0.1	1	2	10
training set	RMSE / meV/Å					
(a) 2000× ICF		63.3	63.2	63.0	63.6	69.5
(b) 20× 100× L		71.5	71.0	69.2	68.9	71.9
(c) 20× $\langle L \rangle_{100}$	71.5	71.0	69.2	71.9		
(d) 20× MF	60.7	60.6	61.4			

Table 5.7: Comparison of GAP performance for different sources of training data. Force prediction RMSE for different values of the noise level hyperparameter σ_f, for GAP trained on (a) ICF for 2000 different configurations, (b) 2000 samples of the constraint force L [20 different CG configurations with 100 samples each], (c) noisy mean forces from the average $\langle L \rangle_{100}$ over the 100 constraint force samples for each of the 20 CG configurations and, for reference/comparison, (d) converged mean forces (MF) corresponding to an average over 2×10^5 constraint force samples.

*The error of the average over 100 samples is smaller than the noise in the individual samples by a factor of 10.

5.5. Summary

In this chapter we applied the GAP-CG approach to small clusters of methanol and to bulk systems of methanol and benzene. We showed for dimers and trimers that the ensembles of the GAP-CG MD simulations are indistinguishable from the ensembles of the coarse-grained coordinates in the atomistic ensemble. For larger systems we did not always achieve a perfect reproduction of the atomistic distributions, but GAP-CG still provided much more accurate results than pair potentials. Even though structure-matching approaches such as IBI reproduce pairwise distance distributions/RDFs by construction, we demonstrated that this does not, in general, allow them to capture more complex properties such as ADFs or correlations involving the orientation of molecules.

We demonstrated that in bulk simulations of the liquid phase a large contribution to the structure comes from environmentally-mediated contributions to the effective interactions due to the relatively close packing: simulations using only a very short-ranged repulsive core potential already reproduce the all-atom structure to a large extent. Hence, in bulk models the lack of flexibility of pair potentials is not as prominent.

Force-based approaches to coarse-graining do not necessarily provide a good reproduction of structural properties. This is indeed the case when only considering pairwise interactions between sites. Therefore, accurately reproducing the structure using a force-based approach is a powerful demonstration of having achieved a consistent CG simulation. Whilst our GAP-CG approach is based on reproducing the mean forces, we can nevertheless recover structural properties of a CG system just as well, if not better than, the more common structure-based approaches that rely on pair potentials.* Importantly, we showed that the GAP-CG potentials capture higher-order correlations very well, which pair potentials fail to do. This suggests that GAP-CG provides a bridge between force-based approaches (such as MSCG/FM) that aim to obtain correct thermodynamic behaviour but do not necessarily recover the structure and structure-based approaches (such as IBI) that obtain correct RDFs but do not reproduce thermodynamic properties.

Accurately reproducing distributions of the all-atom ensemble in the GAP-CG models implies that we recovered good approximations to the respective PMFs. This supports our assumption of separating the PMF into monomer, dimer, and trimer contributions and suggests that GAP will be generally applicable to a wide range of systems.

*In the limit of arbitrarily flexible many-body interactions, both force-based and structure-based approaches will capture the PMF exactly [84].

For the cases presented here, the GAP-CG simulations were slower than even the all-atom simulations. Each GAP force evaluation is significantly more expensive than the evaluation of a pairwise spline interaction, and in the systems we studied in this chapter there was not a sufficient reduction in the number of interactions compared to the all-atom model to make up for the increased computational cost of each interaction.

GAP-CG will only be faster than all-atom simulations when the number of CG interactions is much smaller than the number of atomic interactions. This will be the case for solvated systems with a large number of solvent molecules when we are only interested in the behaviour of the solute. In the next chapter, we will see for the example of benzene solvated by water that the accuracy of GAP allows us to generate a solvent-free CG model that reproduces the properties of the all-atom model significantly more accurately than alternative CG approaches while also providing a significant speedup over all-atom simulations.

6. Practical Application: Solvent-Free Coarse-Graining

In the previous chapter, we showed that we can recover the many-body PMF to high accuracy by using GAP to learn local contributions to the free energy from mean or instantaneous collective forces. While this demonstrated the *accuracy* of the GAP-CG approach, for the systems discussed so far GAP would not be useful in practice, as it is actually *slower* than simulations of the underlying all-atom model: each calculation of a GAP interaction between two molecules is significantly more expensive than the cost of computing the sum of the corresponding molecular mechanics interactions in the underlying all-atom model.

GAP for coarse-graining will only be of practical benefit when it is significantly faster than the underlying all-atom MD. This can be achieved when the CG model has far fewer sites and hence interactions that need to be calculated than the all-atom model on which it is based. A prime example where this is realised are solvent-free CG models for systems which contain a large volume of solvent molecules. This is common in biomolecular simulations where we need to include all the water molecules that make up the bulk of every cell but are more interested in the behaviour of the solutes, such as nucleic acids, proteins or lipids. By discarding all the water molecules we can obtain significant computational gains using GAP models, while the molecular descriptors can represent the effective interactions between solutes, mediated by the solvent molecules, much more accurately than the commonly used site-based potentials.

As an example, we here consider benzene solvated by water. Benzene is hydrophobic, which leads to interesting solvation effects. We describe the all-atom model and discuss the hydrophobic effect in Section 6.1.1. As discussed in Section 5.3, many of the structural properties of bulk systems arise purely due to packing and other environmental effects. In Section 6.1.2 we highlight the influence of the surrounding environment on pair interactions by comparing structural properties and the COM PMF for pairs of benzene molecules in bulk, in gas phase, and dissolved in water at two different concentrations. In a solvent-free CG model, the structural effects that arise

due to the removed solvent molecules need to be introduced by the CG interactions. In Section 6.1.3 we discuss different CG potentials for a solvent-free CG description of water-solvated benzene. In Section 6.1.4 we present results from MD simulations using the CG potentials. We will see that pair potentials cannot satisfactorily account for the structural effects due to the removed water molecules, even if the potentials are constructed to reproduce the solute–solute RDF of the all-atom model. In contrast, we will see that the GAP approach reproduces structural properties much more accurately. A further demonstration of the accuracy of the GAP approach is the reproduction of the COM PMF, which we discuss in Section 6.2. Finally, in Section 6.3 we compare the computational cost of the different approaches.

6.1. Water-Solvated Benzene

6.1.1. All-Atom Description

The structure of molecular benzene was described in Section 5.1. Here we study benzene solvated by water. To describe the structure and interactions of the water molecules in the all-atom simulations we use the TIP3P three-site water model [136].

We consider two different systems of water-solvated benzene: a benzene dimer solvated by 874 water molecules and an octamer of eight benzene molecules solvated by 868 water molecules. In each case, we use a cubic simulation box with periodic boundary conditions and follow the equilibration protocol given in Table 5.5. This results in box lengths of $L_2 = 29.547\,\text{Å}$ and $L_8 = 29.804\,\text{Å}$, respectively, where the subscripts indicate the number of benzene molecules in the system. Hence, the benzene concentrations are $c_2 = 129\,\text{mol}/\text{m}^3 \approx 0.13\,\text{M}$ and $c_8 = 502\,\text{mol}/\text{m}^3 \approx 0.5\,\text{M}$.

The benzene molecules are hydrophobic, which leads to an additional, effective attraction between the solvated molecules compared to bulk benzene. This is illustrated in Fig. 6.1, where we show the distribution of water molecules between two benzene molecules at a range of COM distances, juxtaposed with the COM PMF and the mean force.

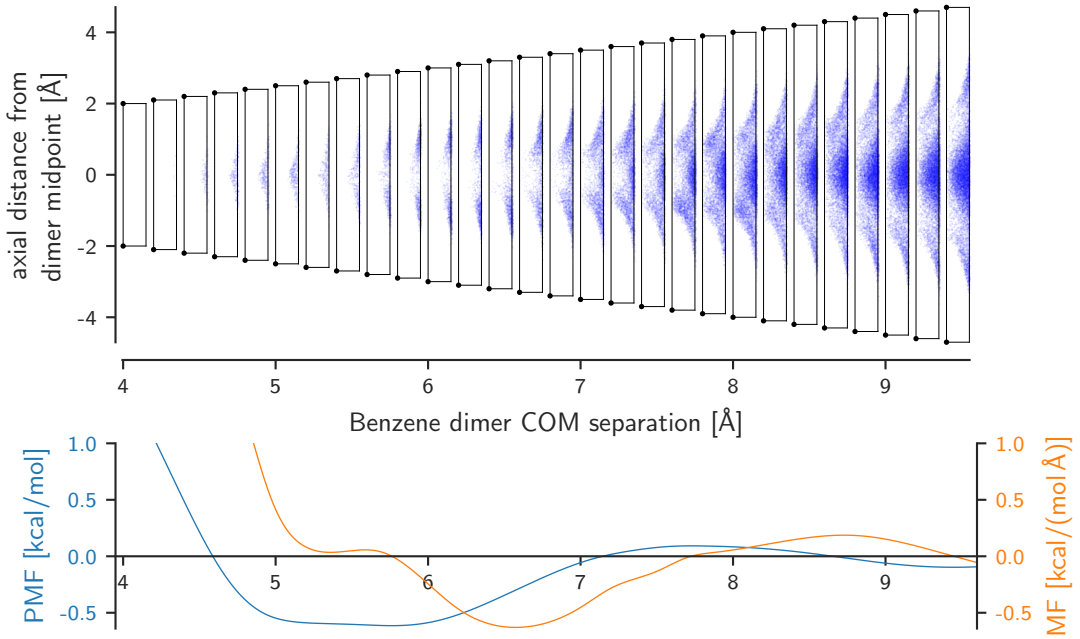

Figure 6.1: Benzene dimer in water: distribution of water molecules.
For a range of COM distances between the benzene molecules, the distribution of the water molecules is shown in blue, measured in terms of axial and radial distance of the oxygen atom from the line connecting the benzene COMs. Each box corresponds to a radial view of a cylinder with radius 2.5 Å; the black circles indicate the positions of the benzene COMs.
Up to a COM separation of 7 Å the water molecules only encroach the edges of the cylinder between the two benzene molecules. Beyond 8 Å there is always a water molecule between them. The "sharing" of the hydrophobic pocket with no water between two benzene molecules effectively gives them an additional attraction. This is reflected by the behaviour of mean force (MF) and PMF for the benzene COM separation.

6.1.2. Benzene Dimer in Different Environments

The interactions between molecules depend significantly on their environment. We demonstrate this at the example of a benzene dimer in gas phase, in bulk benzene, and in bulk water.

In each environment we carried out mean force calculations in which we constrain the centres of mass of two benzene molecules at a range of fixed separations; details are given in Appendix A.4. The resulting COM PMFs are shown together with the RDFs in Fig. 6.2.

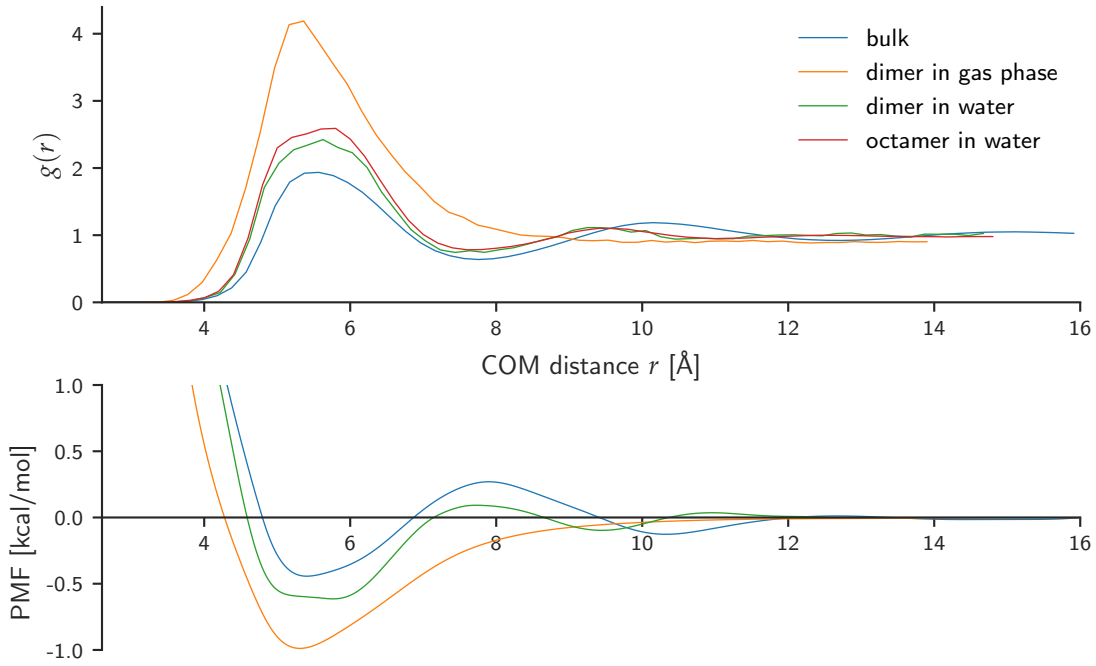

Figure 6.2: Centre-of-mass RDF and PMF for benzene dimers in different environments. We compare the RDF (top) and the PMF (bottom) for bulk benzene, a benzene dimer in gas phase ($V_m = 6.6\,\mathrm{dm}^3/\mathrm{mol}$), a benzene dimer solvated by water ($c_2 \approx 0.13\,\mathrm{M}$), and water-solvated benzene at a higher concentration of $c_8 \approx 0.5\,\mathrm{M}$. The steeper slope in the PMF for the dimer in water corresponds to a stronger mean force due to the hydrophobic attraction.

We also compare the orientational and distance distributions of a pair of benzene molecules in the different environments in Figs. 6.3 and 6.4 on the following pages. All simulations were carried out in a cubic simulation box with periodic boundary conditions. The bulk benzene results are from Section 5.3.4. The gas phase results are for a dimer in a periodic box with $L = 28\,\text{Å}$, corresponding to a molar volume

of $V_m = 6.6\,\mathrm{dm}^3/\mathrm{mol}$. The results for water-solvated benzene are for the systems introduced at the beginning of this section.

It can be seen that structural distributions of a benzene dimer in water are very different from those in gas phase. In fact, the distributions of the water-solvated benzene are more similar to those in bulk benzene. The RDF shows an oscillatory structure with maxima where an integer number of water molecules neatly fit in the space between the benzene molecules. We also note that the distributions for water-solvated benzene dimer and octamer closely resemble each other; this suggests that the benzene concentrations are too small for (molecular) three-body effects to play a significant role.

In the remainder of this chapter we will see that pair potentials between CG sites do not allow enough flexibility to represent the differences between effective interactions in the different environments, whereas these differences can be captured by the molecular interaction terms available in the GAP approach.

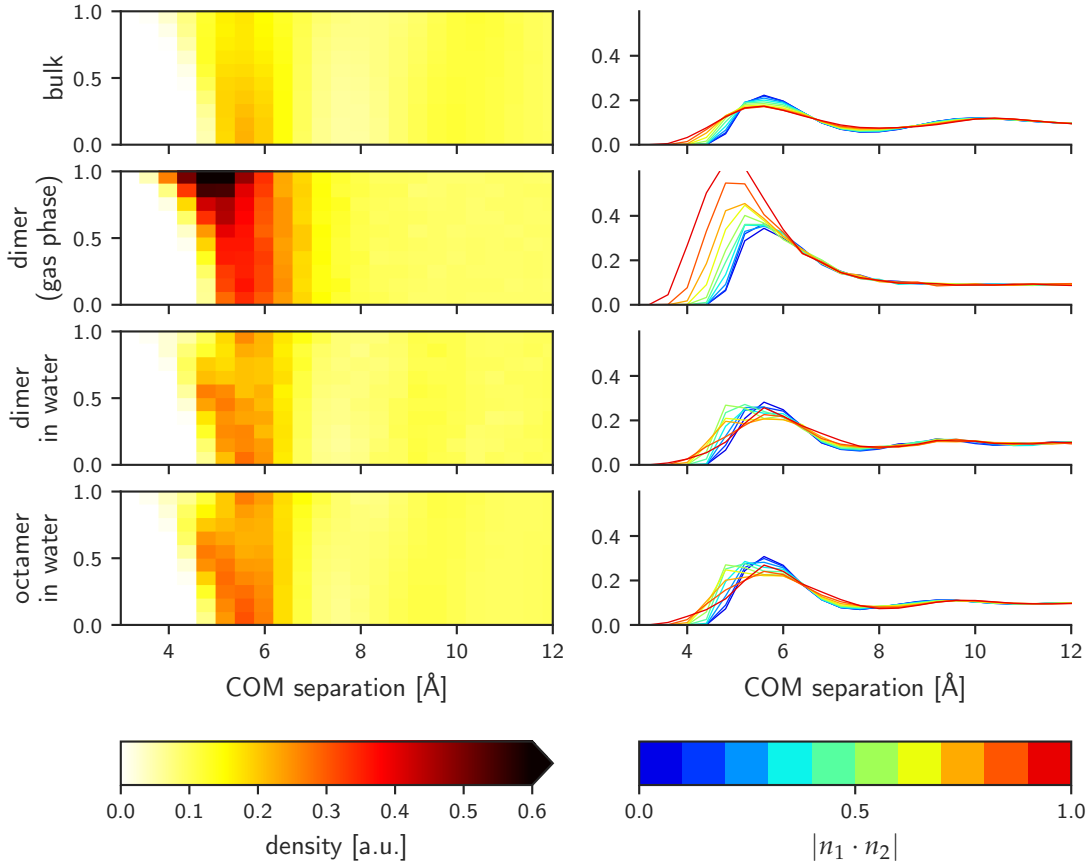

Figure 6.3: Orientational distributions of benzene pairs in different environments (1/2). We show the distribution of plane orientations, measured by the absolute value of the dot product between the plane normal vectors, $|n_1 \cdot n_2|$, as a function of the COM separation. *Left:* 2D distribution. *Right:* horizontal slices through this distribution; each line corresponds to a certain range of dot product values. The distributions have been normalised according to the RDF convention along the distance axis; the sum over orientations corresponds to the COM RDF. Uniformly random orientations would lead to a uniform profile of the dot product, corresponding to all lines lying on top of each other.
Note that the pair structures of the two water-solvated systems are very similar to each other, and noticeably different from bulk benzene. However, the differences between a pair of benzene molecules in bulk water and in bulk benzene are small compared to the much larger difference to the gas phase dimer.

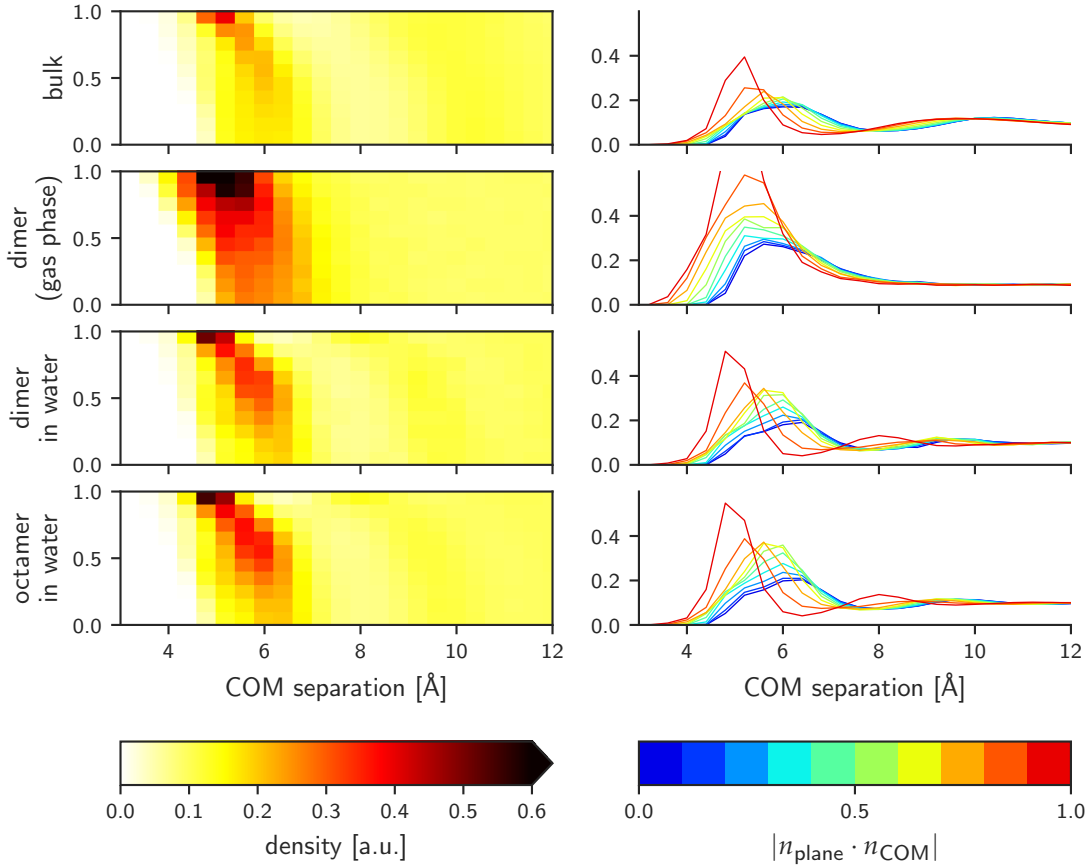

Figure 6.4: Orientational distributions of benzene pairs in different environments (2/2). This figure is similar to Fig. 6.3; here, we measure the orientation using the absolute value of the dot product between the plane normal vector n_{plane} of one benzene molecule and the normal vector n_{COM} in the direction of the COM separation. Again, there is a strong similarity between the pair structures of water-solvated benzene at different concentrations, clearly distinct from the bulk benzene structure, and the gas phase structure is very different from the structures in bulk.

6.1.3. CG Potentials

As before, we compare force-based and structure-based approaches to coarse-graining. For the force-based approaches, we collected mean forces for 17000 different CG configurations of the water-solvated benzene dimer in constrained MD runs of 100 ps. From these mean force observations we trained a GAP with the molecular dimer descriptor and generated a force-matching pair potential using the MSCG/FM method. Structure-based approaches are represented by the pair potential obtained by the IBI approach for the water-solvated benzene octamer. For comparison, we also include a "traditional" implicit solvent model of the solvated benzene dimer in which the solute is described in atomistic detail, based on the Generalised Born model for solvation effects [137].

GAP

The hyperparameters we used were similar to bulk benzene, *cf.* Table A.3. On a separate test set of mean forces for 3000 CG configurations, this potential has an RMSE of 17 meV/Å.

Based on a one-site CG model for benzene in water, Ref. 138 suggests that potentials based on infinite dilution can transfer well to larger concentrations. Motivated by this, we take the dimer GAP trained entirely on dimer configurations and also apply it at higher concentration in the system of eight benzene molecules in water.

Pair potentials

The CG description of benzene used in this work has only one site type and, hence, there is just a single non-bonded pair potential. This makes it easy to provide the pair potential-based approaches with sufficient data to ensure that we have obtained the best effective pair potentials, according to the respective approach. However, it also means that it is difficult for the pair potentials to approximate the true interactions as given by the PMF.

Force-based approach. For the water-solvated benzene dimer, we construct a force-matching pair potential using the MSCG/FM method. The pair potential has an RMSE of 31 meV/Å on the mean force test set. We will see that the force-matching pair potential fails to accurately describe the distributions of the all-atom dimer model. We do not expect a force-matching pair potential to perform significantly better for

the water-solvated benzene octamer, and for this reason we do not consider the force-matching pair potential again when discussing benzene at a higher concentration in water.

Structure-based approach. For the water-solvated benzene octamer we determine a pair potential that reproduces the site–site RDF using the IBI approach as discussed previously. On the test set of mean forces in the solvated dimer system, the IBI pair potential shows an RMSE of 36 meV/Å. We do not include a structure-matching pair potential for the water-solvated benzene dimer, as it would have been very computationally expensive to generate sufficiently converged RDFs at lower concentration.

The resulting pair potentials and the pair forces are shown in Fig. A.5 in the Appendix. We compare the force predictions of the CG potentials along cutlines through the configuration space of the benzene dimer in Fig. 6.5. This shows that the dimer GAP reproduces the all-atom mean forces significantly more accurately than either of the pair potentials, though we note that in this case there are some visible discrepancies between the all-atom reference and the GAP predictions.

Atomistic implicit solvent model

For comparison, we include results from simulations using a "traditional" implicit solvent model, in which the solute molecules are modelled with atomic detail. In such models, the solvation effect is described implicitly by changing the force field to account for electrostatic effects that are due to the bulk solvent. As we will see, this does not reproduce structural effects that are due to the absent solvent molecules. We use the Generalised Born solvation model implemented in the AMBER code [139, 140]. This type of model requires vacuum boundary conditions; hence, we only applied it to the water-solvated benzene dimer.

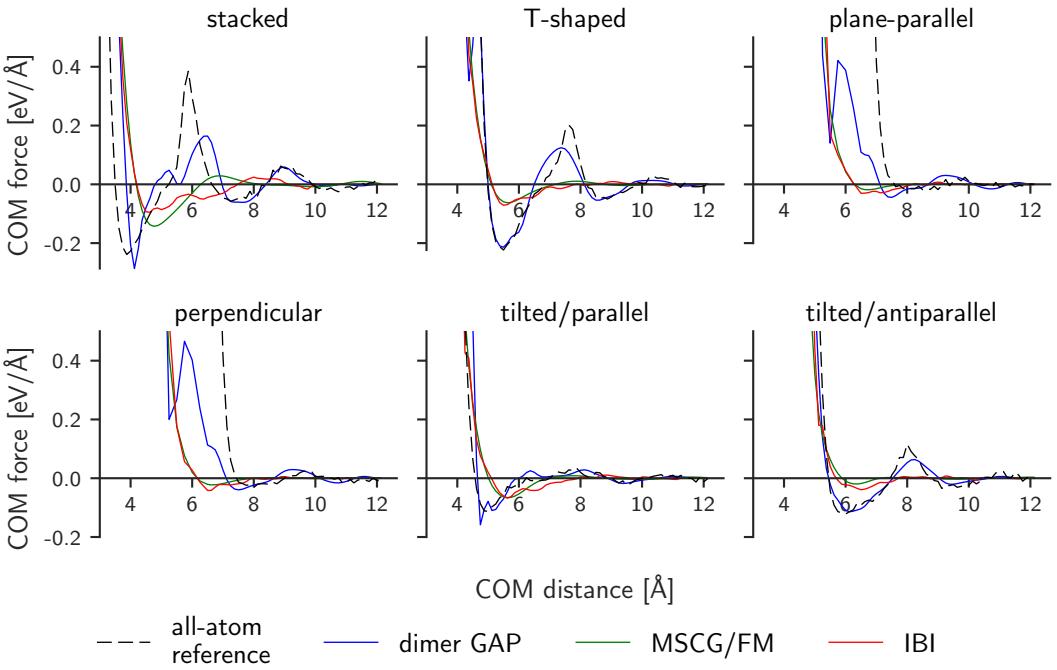

Figure 6.5: Water-solvated benzene dimer: COM force along cutlines.
This shows the mean force along the COM separation for six different relative orientations of two benzene molecules solvated by water as a function of the COM distance, compared to the force predictions by the dimer GAP potential and the force- and structure-matching pair potentials. Though none of the potentials reproduce the mean forces exactly, the dimer GAP is the only one that qualitatively captures all the different orientations.

6.1.4. MD Results

Water-solvated benzene dimer

We compare the all-atom model with the dimer-descriptor GAP, the force-matching pair potential and the atomistic implicit solvent model. As was the case for the methanol clusters, we impose a half-open harmonic restraint on the dimer COM separation with a force constant of 50 kcal/mol/Å^2, beginning at a separation of 12 Å. The resulting distance distributions are shown in Fig. 6.6, and results for orientational distributions are shown in Figs. 6.8 and 6.9. It can be seen that the atomistic implicit solvent model is a better approximation to the dimer solvated by explicit water than a pure gas phase simulation, and marginally better than the force-matching pair potential. However, both the atomistic implicit solvent model and the pair potential fail to reproduce RDF and orientational distributions correctly. In contrast, the dimer GAP model reproduces the RDF and even the orientational distributions. Whereas the distribution of $n_{\text{plane}} \cdot n_{\text{COM}}$ is reproduced almost perfectly, there is a clearly visible discrepancy in the distribution of $n_1 \cdot n_2$.

Water-solvated benzene octamer

We compare the all-atom model with the dimer-descriptor GAP trained on dimer MF and the structure-matching pair potential based on reproducing the site–site RDF. The COM and site–site RDFs are shown in Fig. 6.7, and the COM ADF is shown in Fig. 6.12. The orientational distributions are shown in Figs. 6.10 and 6.11. Whereas the structure-matching IBI pair potential reproduces the orientational distributions somewhat better than the force-matching pair potential, it does not reproduce the oscillations of the COM RDF correctly and provides less accurate results than the dimer GAP. Moreover, the discrepancy between the dimer GAP ADF and the all-atom ADF is on the same scale as the difference between the all-atom RDFs for different concentrations; hence, a dimer GAP trained on configurations of the water-solvated *octamer* would be expected to reproduce the structural distributions even more closely.

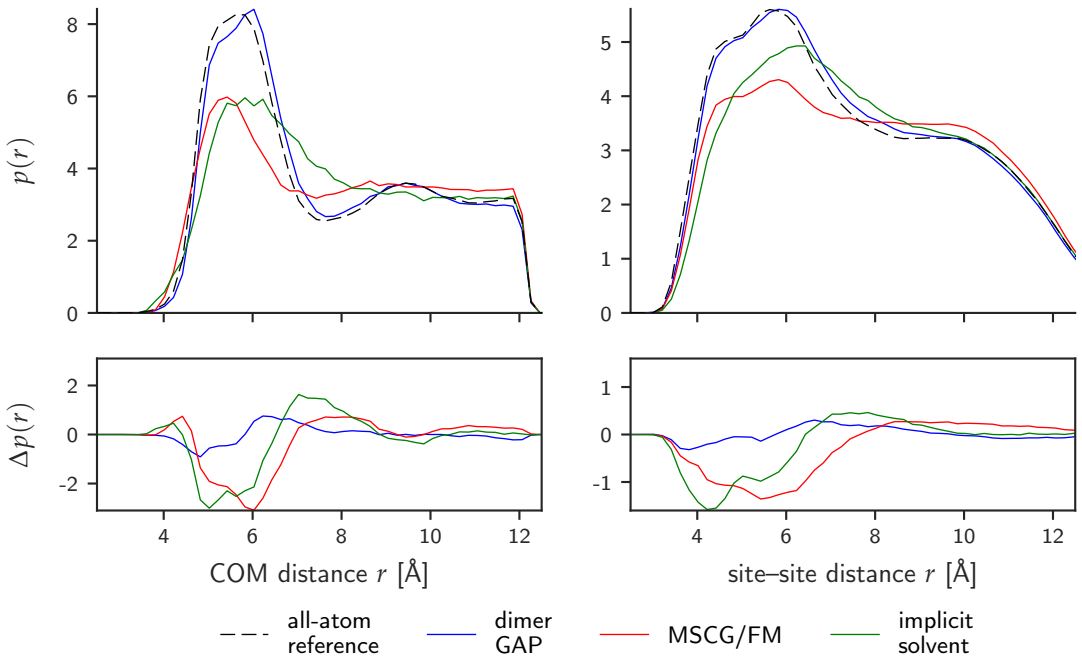

Figure 6.6: Water-solvated benzene dimer: distance distributions.
We compare the dimer GAP, the MSCG/FM pair potential, and the atomistic implicit solvent model with the all-atom reference model that explicitly includes water. In each case, a half-harmonic restraint was applied to the COM distance starting at 12 Å, with force constant 50 kcal/(mol Å^2). Only the many-body dimer GAP is able to capture the effective interactions and reproduce the all-atom distance distributions. Note that the water molecules in the all-atom simulation mediate an effective interaction between the benzene molecules even beyond the all-atom force field cutoff of 9 Å.

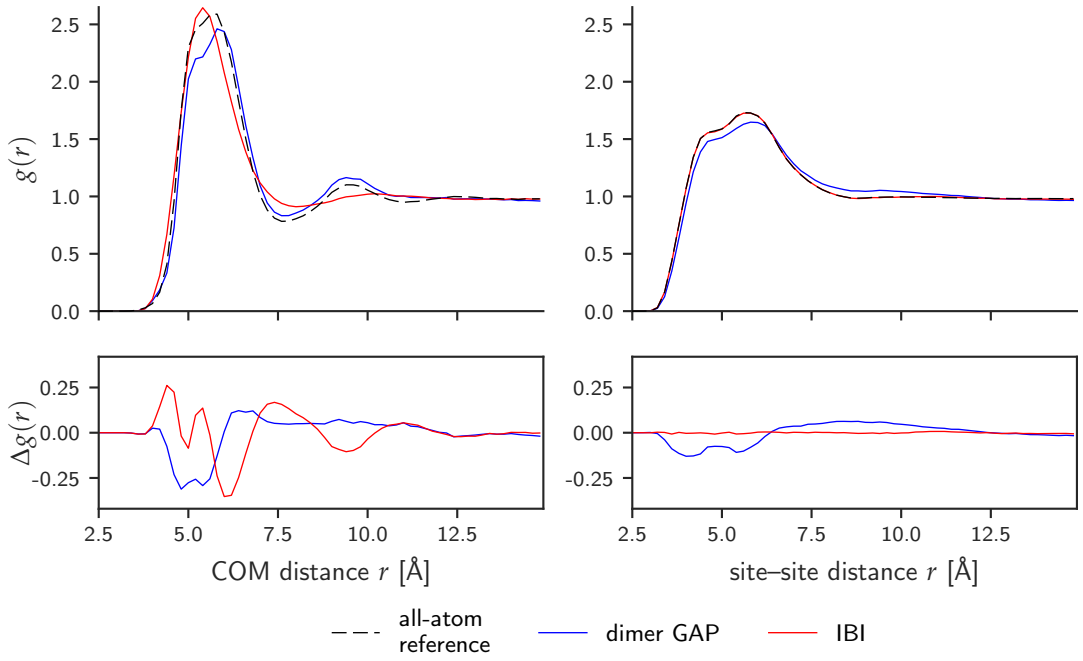

Figure 6.7: Water-solvated benzene octamer: COM and site–site RDF.
By construction, the IBI pair potential reproduces the site–site RDF exactly. Nevertheless, it fails to capture the oscillatory structure in the COM RDF. In contrast, the dimer-descriptor GAP reproduces the RDFs to a reasonable accuracy, even though it was trained on water-solvated *dimer* configurations.

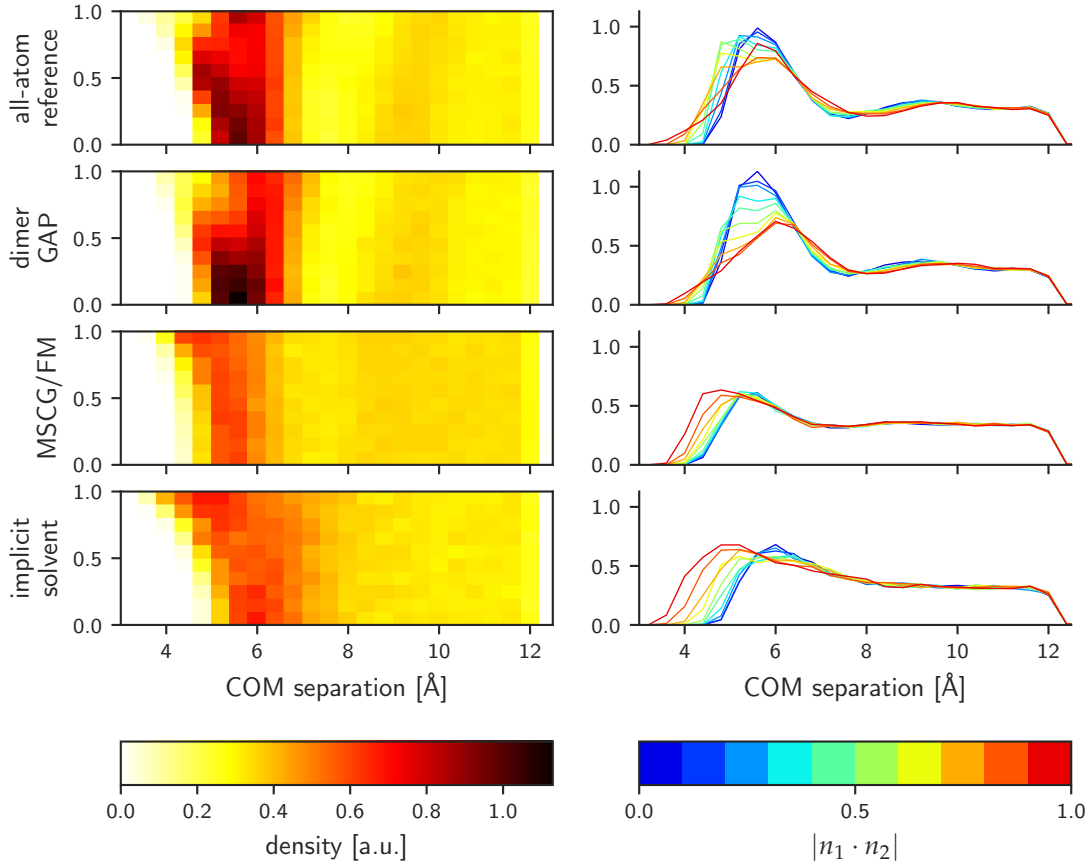

Figure 6.8: Water-solvated benzene dimer: orientational distributions (1/2). Plane orientation measured by the absolute value of the dot product between the plane normal vectors, $|n_1 \cdot n_2|$; for more explanation see Fig. 6.3. This shows that the dimer GAP provides a much better approximation to the all-atom model with explicit water than MSCG/FM pair potential or atomistic implicit solvent model.

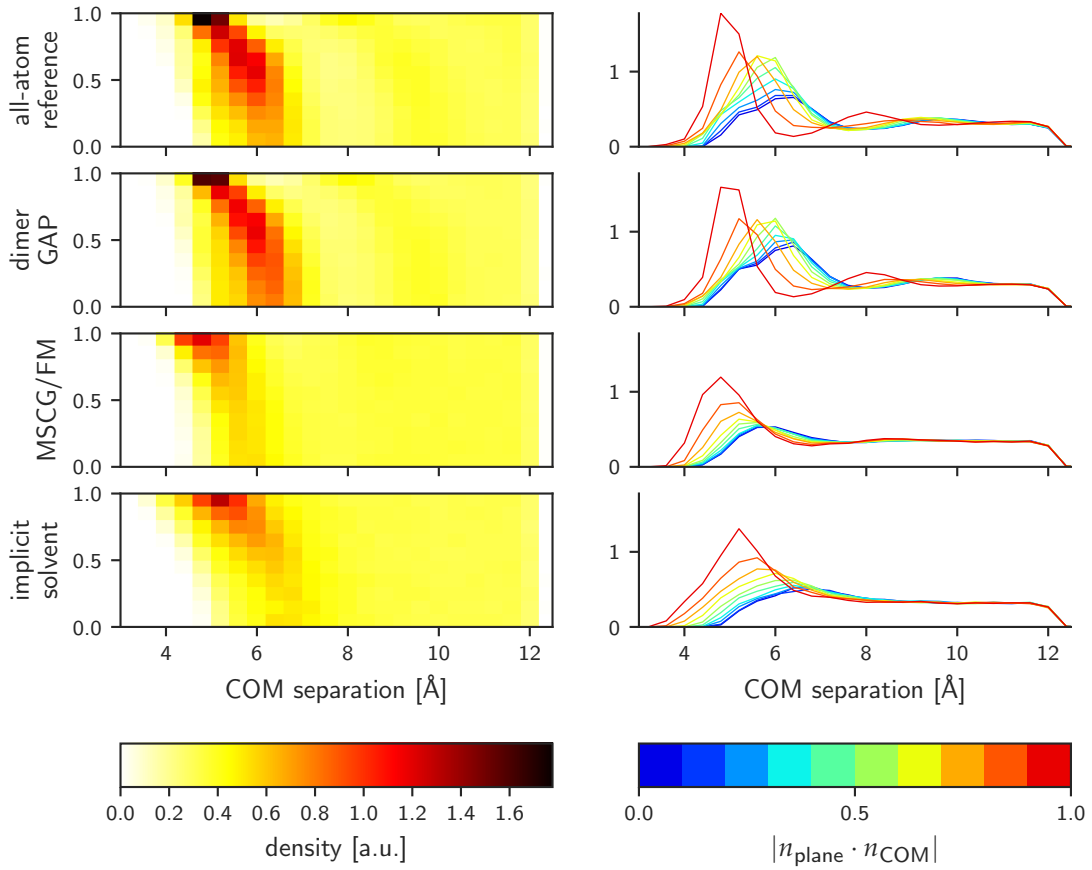

Figure 6.9: Water-solvated benzene dimer: orientational distributions (2/2). Plane orientation measured by the absolute value of the dot product between the plane normal vector n_{plane} of a benzene molecule and the direction n_{COM} between the benzene COMs. The all-atom reference distribution is accurately reproduced by the dimer GAP, whereas MSCG/FM pair potential and atomistic implicit solvent model show significant discrepancies.

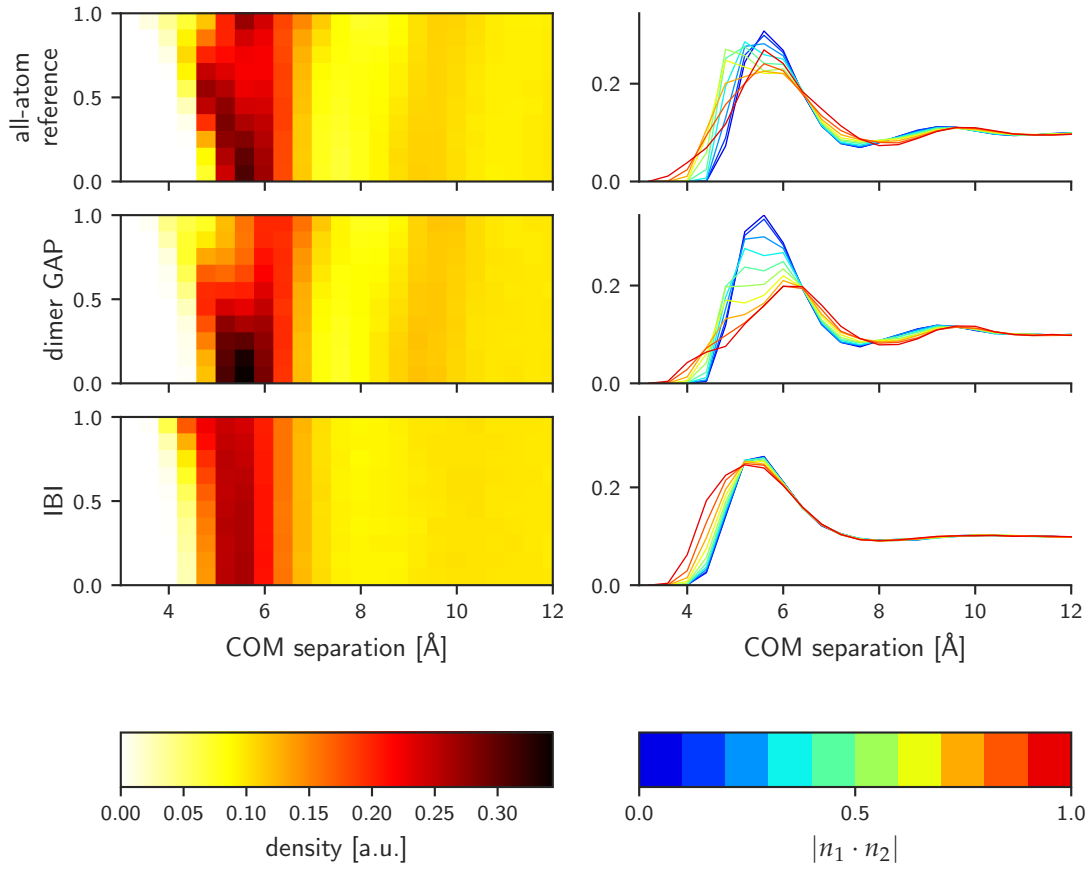

Figure 6.10: Water-solvated benzene octamer: orientational distributions (1/2). Dot product between plane normal vectors against the COM separation for all pairs of benzene molecules.

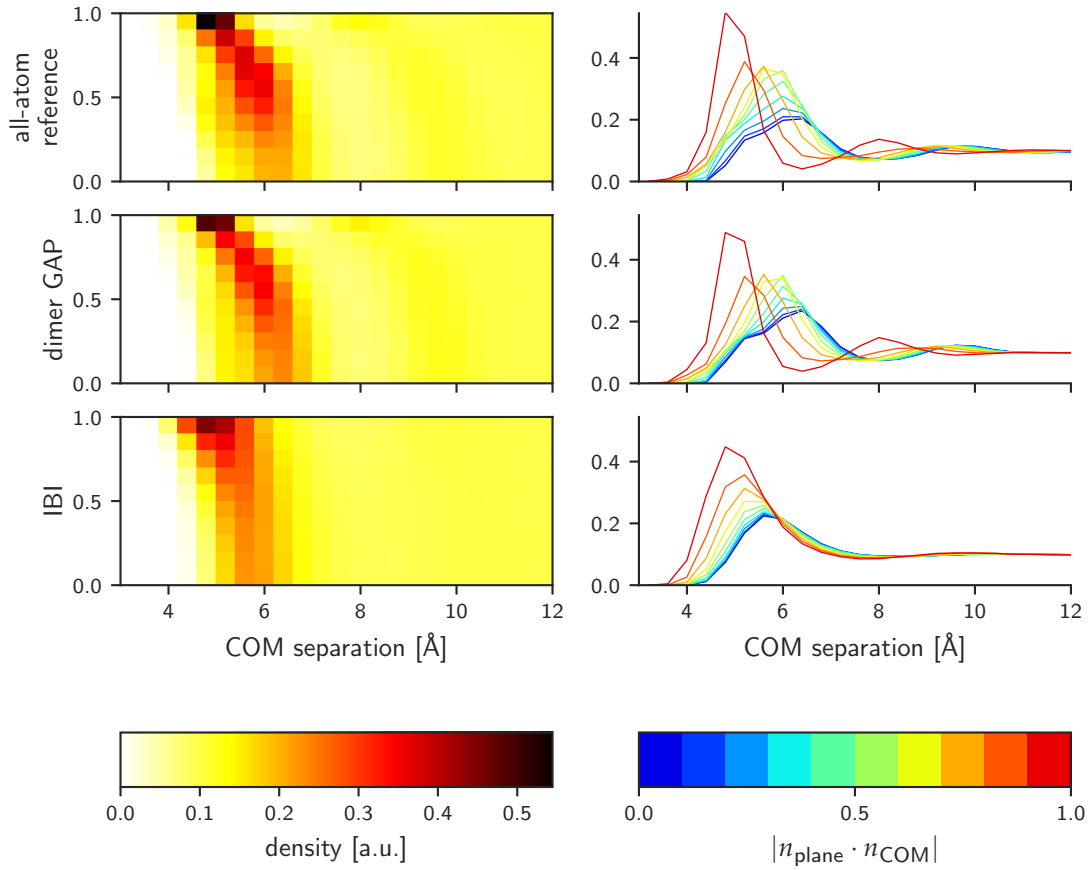

Figure 6.11: Water-solvated benzene octamer: orientational distributions (2/2). Dot product between plane normal vector and normal vector between COMs against the COM separation for all pairs of benzene molecules.

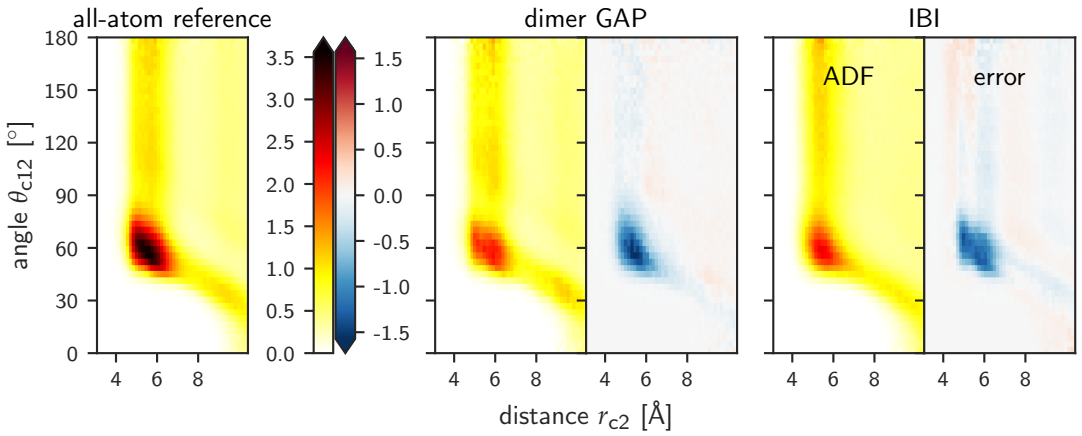

Figure 6.12: Water-solvated benzene octamer: Angular Distribution Function of the benzene COMs, given by the 2D distribution of the angle θ_{c12} between the centres of mass of a central molecule and two neighbours against the distance r_{c2} to the second neighbour. Only triplets in which the distance r_{c1} to the first neighbour is less than 7 Å were included. Note that the difference between dimer GAP (trained on dimer MF) and all-atom reference ADF is in line with the difference between the all-atom RDFs for water-solvated benzene dimer and octamer, *cf.* Fig. 6.2 on page 156.

6.2. Recovering a Pair PMF from the CG Potential

As a further demonstration of the power of the GAP approach, we compare the GAP-predicted PMF with the COM PMF obtained from Constrained MD with constraints on the centres of mass of two benzene molecules in water (as shown in Fig. 6.2).

GAP predicts an approximation to the N-body PMF $A(\boldsymbol{R}^N)$ for all CG coordinates after averaging over the "internal" atomistic degrees of freedom. To compare this with the one-dimensional COM PMF $A_{\mathrm{COM}}(r)$ shown in Fig. 6.2, we need to further average over the CG degrees of freedom orthogonal to the COM distance r between the two molecules, corresponding to Eqs. (2.4.2) and (3.3.7) on a higher level:

$$A_{\mathrm{COM}}(r) = \int \mathrm{d}\boldsymbol{R}^N \, \exp\big(-\beta A(\boldsymbol{R}^N)\big)\delta\big(r - r(\boldsymbol{R}^N)\big).$$

For the water-solvated benzene dimer, this procedure corresponds to a Boltzmann average of the PMF over rotations and deformations of the monomers. Fortunately, within the GAP approach, contributions to the PMF have already been separated into monomer and dimer contributions. The monomer descriptor captures the contributions that are independent of the presence of other molecules and, hence, the distance between the two molecules. Therefore, we can effectively obtain the Boltzmann average over all monomer configurations by simply discarding the monomer prediction and only considering the contribution to the PMF predicted by the dimer descriptor.

We carry out the Boltzmann average over rotations by generating 1000 random, uniformly distributed orientations for the benzene monomers and predicting the dimer contribution for these orientations for a range of centre-of-mass separations. The results are shown in Fig. 6.13 on the following page, where it can be seen that the dimer GAP prediction closely follows the COM PMF obtained in all-atom MD, whereas the pair potentials fail to describe its oscillatory structure.

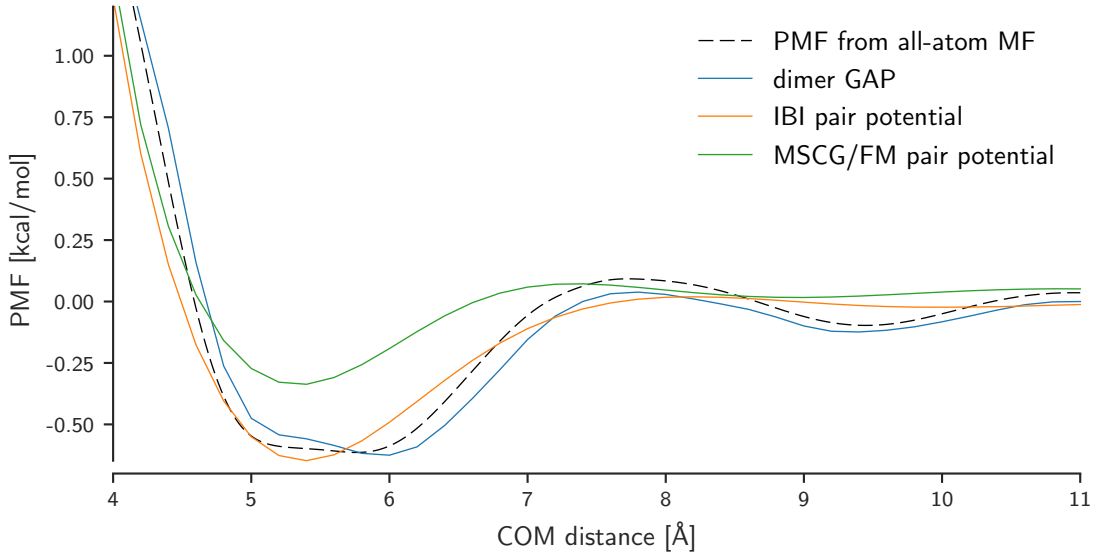

Figure 6.13: COM Potential of Mean Force. We compare the PMF obtained through all-atom Constrained MD simulations with the COM PMF predictions of the CG interaction potentials. The COM PMF predictions were obtained from calculating the Boltzmann-weighted average over monomer orientations of the dimer terms in the potentials. This supports our claim that the GAP approach can accurately capture the many-body Potential of Mean Force and separate it into monomer and higher-order contributions.

6.3. Simulation Speed

The overall speed of a molecular dynamics simulation for a given model or interaction potential is defined in terms of how much model time we can simulate in a certain amount of wall-clock time, typically measured in ns/day. This depends on the system size, the number of computing cores, and the time step Δt. The run time of molecular dynamics simulations generally scales linearly* with the number of particles N. Ideally, doubling the number of cores would halve the time needed to run a simulation. However, due to parallel overheads, the run time will not necessarily scale inversely with the number of cores. In Table 6.1 we compare the speed of all-atom simulations with that of CG pair potentials and dimer GAP. Parallel scaling behaviour of GAP simulations compared to all-atom MD is shown in Fig. 6.14.

*The Ewald summation that deals with electrostatic interactions in periodic systems actually scales as $\mathcal{O}(N \log N)$, but the difference to linear scaling is negligible in practice.

system	model	Δt / fs	N_{site}	cores	ms/step	ns/day	$\frac{\text{ns/day}}{\text{core}}$
dimer	all-atom	0.5	2646	16	2.4	18	1.1
octamer	all-atom		2700	16	2.4	18	1.1
				8	3.9	11	1.4
				4	7.2	6	1.5
				2	14	3	1.5
				1	28	1.5	1.5
dimer	dimer GAP	4	6	1	44	7.9	7.9
octamer	dimer GAP		24	4	42	8.2	2.0
octamer	IBI pair potential		24	4	0.14	2400	600

Table 6.1: Simulation speeds for water-solvated benzene dimer and octamer.

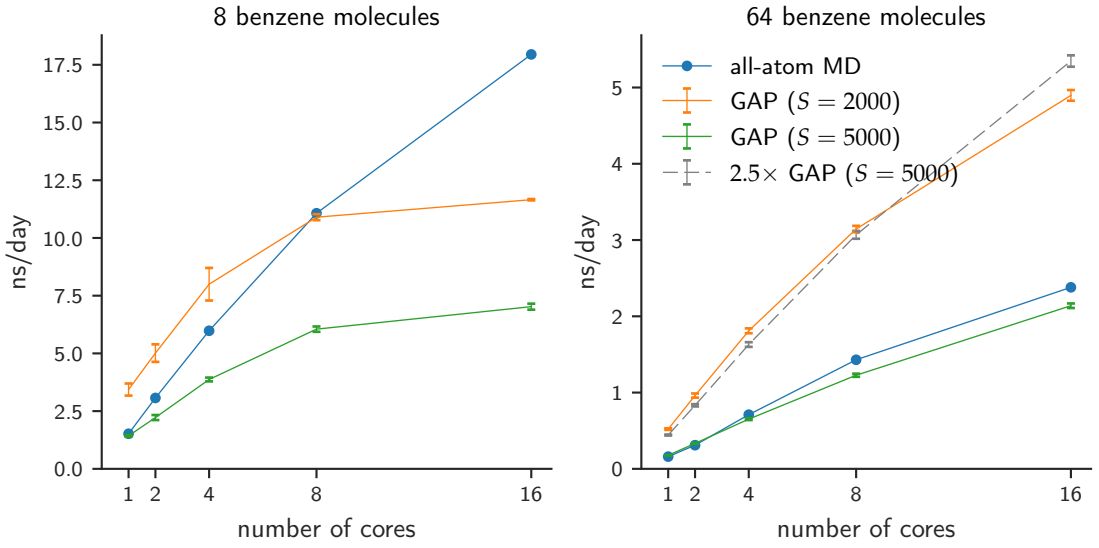

Figure 6.14: Parallel scaling. Simulation speeds for all-atom and dimer GAP MD, where S indicates the number of sparse points. Ideal scaling would correspond to a linear relationship between simulation speed and number of computing cores. These results show that the parallel scaling improves for larger system sizes. The dashed line shows the result for $S = 5000$ sparse points scaled by the factor 2.5; this demonstrates the linear scaling with S, and also underlines that larger calculations scale better to a higher number of cores.

6.4. Summary

Here we have demonstrated a GAP-CG model that is significantly faster than the underlying all-atom simulation, while providing a good reproduction of the all-atom distributions and correlations. In contrast, implicit solvent simulations using pair potentials, whether describing the solute with atomic detail or using CG sites, fail to reproduce all-atom distributions and correlations.

We note that the water-solvated benzene is expected to be at the lower end of speedups that are attainable in solvent-free CG models using GAP. Here, the time step in the CG simulations is limited by the comparatively stiff monomer term. When considering a more drastic coarse-graining of larger molecules, the resulting monomer interactions are likely to be softer, thereby allowing a larger time step in the CG simulation. Moreover, larger molecules would require a larger box, resulting in an even higher speedup compared to all-atom simulations. In cases where real dynamics are not relevant, the CG simulations are further accelerated by the smoother CG interactions, which lead to reduced friction in the CG model.

We acknowledge that the GAP for the water dimer did not reproduce all of the all-atom distributions perfectly. There is also a visible discrepancy between GAP-predicted and all-atom mean forces along parts of the cutlines. However, there are no intrinsic reasons that would prevent the dimer GAP from accurately reproducing the water dimer PMF. Investigating this discrepancy will be a subject of future research.

In further work, this GAP-CG model could, for example, be used in a random search for minima. This would not be feasible with the all-atom potential due to the numerous local minima of the solvent configuration in which any minimiser would get stuck.

7. Conclusions

In this thesis I showed that the Gaussian Approximation Potential (GAP), previously used for interatomic potentials, can recover the many-body potential of mean force of a coarse-grained system to a high accuracy. I demonstrated for the examples of methanol clusters, bulk methanol and benzene, and water-solvated benzene clusters that molecular dynamics simulations using this GAP-CG approach are significantly better at reproducing distributions and correlations of the underlying fine-grained model than currently predominating approaches. Whereas other CG approaches generally describe effective interactions in the coarse-grained model using pairwise radial interactions between CG sites, GAP is based on molecular terms that assign contributions to the free energy to the internal configurations of monomers and to dimer and trimer configurations. Using Gaussian process regression, GAP infers these free energy contributions from observations of mean or instantaneous collective forces on the CG sites.

Coarse-grained interaction potentials that are based on matching forces were generally thought not to be as good at recovering *structural* distributions as approaches that directly focus on reproducing the structure, for example, by iteratively refining potentials to match Radial Distribution Functions (RDFs). In contrast, such structure-based approaches generally fail to reproduce thermodynamic properties of the underlying model. In the present work I demonstrate that this lack of representability is actually due to the mistaken assumption that pairwise interactions between individual sites can provide a good approximation to the true CG interactions. The flexibility of the monomer, dimer, and trimer descriptors in the GAP approach can represent general many-body interactions (for example, a molecular dimer descriptor for a three-site description of a CG molecule corresponds to 6-body interactions in terms of CG sites). This allows us to recover the structural distributions of the underlying system *better* than structure-based approaches that only optimise the interactions based on a subset of distributions (usually the RDFs). GAP provides a flexible and powerful approach for representing not just interatomic potential energy surfaces, but also, as this thesis has shown, for coarse-grained free energy surfaces.

Compared to "atomistic GAP" for interpolating the potential energy surfaces of all-atom systems, GAP-CG faces additional challenges. Whereas both function values and gradient observations are available for the potential energy surface, in the coarse-graining approach we only have access to the gradients (mean or instantaneous collective forces) and from these need to reconstruct the free energy surface. This makes it harder both to train an accurate GAP and to validate its accuracy. Without function value observations, the Gaussian process regression requires more training data to reconstruct the unknown energy surface to comparable accuracy. Moreover, when the GAP is trained from instantaneous collective forces, the training set contains a significant amount of noise. When mean forces are available, the GAP force predictions can be compared with mean forces for a separate test set of configurations that were not used in the training. Based on comparison with the test set mean forces, it is possible to tune the hyperparameters and thereby optimise the GAP predictions. When accurate mean forces are not available as a reference, we need to rely on intuition for selecting the values of the hyperparameters. However, in the present work I developed working values for the hyperparameters and showed that the GAP prediction accuracy is fairly insensitive to changes of the hyperparameters, so this is not as serious an issue as might have been thought initially.

Many-body interactions are computationally expensive. The evaluation of a GAP interaction is significantly more expensive than the evaluation of a simple tabulated spline (by two to three orders of magnitude). Hence, in simulations of bulk systems where many interactions need to be calculated, GAP-CG is significantly slower than even the all-atom simulations. In bulk systems, packing effects are responsible for a significant part of the overall structure and, for this reason, CG pair potentials already provide reasonable accuracy.

GAP-CG will be useful in cases where it is faster than all-atom MD whilst being significantly more accurate than CG models based on pair potentials. This can be the case in systems with a large volume of solvent, for which most of the computing time in typical all-atom MD simulations is spent on calculating non-bonded interactions between solvent molecules. If we are interested in the effective behaviour of the solutes, we can remove the solvent degrees of freedom. A solvent-free GAP-CG simulation can significantly reduce the computational cost compared to the all-atom simulation. Even though simpler methods based on pair potentials would be faster still, this would be at the cost of a significantly reduced accuracy with respect to the reproduction of properties of the underlying system.

GAP-CG does not supersede previously introduced approaches, but provides another tool whose usefulness depends on the application. GAP may be especially useful when combined with other approaches. In the following, I discuss scaling to larger system sizes (Section 7.1), limitations of the GAP approach to coarse-graining (Section 7.2), and an outlook to the future of this approach (Section 7.3).

7.1. Scaling

GAP-CG will only provide a speed advantage for CG models that contain a significantly smaller number of CG interactions than the all-atom system. The precise number of interactions depends on the specific configuration and can change in each time step, as atoms and molecules move in and out of the interaction cutoffs. For this reason, it is common practice to simply compare the number of interaction *sites*, and the ratio between the total number of atoms and the number of CG sites provides a proxy for measuring potential speedups due to coarse-graining.

In systems of comparable size to the cutoff distances, the number of dimer interactions increases quadratically with the number of sites. In large systems, because of the cutoffs used in both all-atom and CG models, the number of interactions that have to be computed scales linearly with the system size.

In Chapter 6 I discussed the coarse-graining of water-solvated benzene and showed that GAP-CG MD is faster than the corresponding all-atom simulation (though this depends on the number of sparse points used in the Gaussian process regression). There, the ratio between the number of atoms and the number of CG sites is on the order of one hundred. For larger biomolecular simulations, coarse-graining can lead to significantly larger computational gains, as I will discuss in the following.

As an example, let us consider DNA, the carrier of genetic information in all living organisms. In order to provide a large strand of DNA with the freedom to arbitrarily bend and flex, a sufficiently large simulation box and hence increasingly large volumes of solvating water molecules are needed. By excluding the solvent from the coarse-grained description, we vastly reduce the number of particles and thus the number of interactions that need to be evaluated during the simulation. I have quantified this saving by explicitly solvating DNA systems of different sizes in water and counting the total number of atoms. The results are presented in Table 7.1 on the following page. For comparison, numbers for the water-solvated benzene systems and for the HIV-1 capsid model that was mentioned in the Introduction are also included.

Hence, for larger models, GAP-CG may be faster than all-atom MD by an additional

order of magnitude. For models in which internal fluctuations of the CG monomer occur on longer time scales than for the comparatively stiff (high-frequency) methanol or benzene molecules discussed in this thesis, it may also be possible to further increase the time step used in the GAP-CG MD simulation.

solute	atoms (solvated)	CG sites	ratio
benzene (2 molecules, $c = 0.13\,\mathrm{M}$)	3×10^3	6	5×10^2
benzene (8 molecules, $c = 0.5\,\mathrm{M}$)	3×10^3	24	1×10^2
DNA (1 base pair [bp])	3×10^3	6	5×10^2
DNA (linear, 10 bp)	1×10^4	60	2×10^2
DNA (half-circle, 50 bp)	2×10^5	300	7×10^2
DNA (half-circle, 100 bp)	1×10^6	600	2×10^3
HIV-1 capsid	64×10^6	6580	9×10^3

Table 7.1: Scaling of the number of atoms with system size, compared to the number of CG sites. The solvated benzene systems were discussed in Chapter 6. Solvated atom numbers for DNA are a conservative estimate based on a fixed 8 Å solvent margin around each *initial* configuration of the DNA; in practice, the 50 and 100 bp systems would need an even larger simulation box for an accurate simulation. As for benzene, a three-site CG description is assumed for each DNA base. Numbers for the HIV-1 capsid are based on the data in Ref. 11, assuming a description using five CG sites per capsid protein.

7.2. Limitations of GAP-CG

The main limitation of GAP is its computational cost compared to simpler CG approaches. As GAP consists of independent training and prediction stages, their respective computational costs will be discussed separately.

In the training stage, we need to construct and invert the covariance matrix between sparse points to obtain the sparse coefficients. This involves matrix multiplication with the matrix of covariances between sparse and training points. Overall, the GAP training scales as NS^2 in calculation time and as NS in memory, where N is the number of training points and S is the number of sparse points. Due to the large number of points that are needed to accurately determine the interactions, this stage is primarily limited by the available memory. In many cases, the GAP training requires dedicated large memory nodes; however, this effort only needs to be expended once for a CG model.

In the prediction stage, the forces at a "test point" are obtained by simply calculating the dot product of the sparse coefficients with the vector of covariances between test

point and sparse points. This operation is linear in S,* as demonstrated by Fig. 6.14. The predictions are mainly compute-bound.

The performance of GAP-CG also depends on the dimensionality of the descriptors. Clearly, higher-dimensional descriptors will be more expensive. Increasing the descriptor dimensionality directly increases the computational cost of calculating covariances. More significant, however, is that this will also require a larger number of training and sparse points to accurately describe the local free energy contributions in a higher-dimensional space. For very high-dimensional descriptors, it will become difficult to achieve sufficient data coverage of the descriptor space.

In this work I used descriptors based on the list of *all* pairwise distances between the sites involved in monomers, dimers and trimers. As discussed in Section 4.2.2, this generally provides an *overcomplete* description. We may obtain better performance by constructing cleverer, lower-dimensional descriptors for the systems we want to investigate, for example, by only considering a certain subset of the distances.

When comparing the speed of GAP with other approaches, it is important to remember that GAP is essentially a development code and has not yet benefited from as many years of performance optimisation as common MD packages for established potentials such as AMBER or LAMMPS. In the course of this work, I already significantly improved the parallel performance of GAP MD.

As we saw in Fig. 6.14, the parallel scaling with the number of cores is better for larger systems. Descriptors are all independent of each other and only depend on the positions of CG sites within a localised region. There is no reason in principle that would prevent GAP from parallelising to much larger systems. It is worth noting that MD packages are typically designed with all-atom simulations in mind. Solvent-free CG MD will require new algorithms that take advantage of the sparse nature of the CG system; further discussion of this issue can be found in Ref. 141.

7.3. Outlook and Future Work

In this work I demonstrated that GAP-CG can significantly improve the simultaneous representation of both thermodynamic and structural properties of the fine-grained system by CG simulations. Building on this achievement, there are a range of different directions that would be worth pursuing. I will first discuss the use of GAP to describe the monomer term in highly coarse-grained models. I will then discuss the use of GAP

*The computational cost of predicting the variance scales as S^2.

for a better description of coarse-grained interactions in adaptive resolution simulations. Finally, I will present an outlook towards improving not just representability, but also transferability.

Highly coarse-grained models. GAP is very promising for highly coarse-grained models in which CG sites do not represent just a few atoms, but whole side chains or even whole residues; for example, this can be a lipid of 144 atoms that is represented by only three to five sites [88]. This results in significantly fewer CG interactions, and hence the overhead of GAP might not be too much of a disadvantage. At the same time, potentials of mean force for highly coarse-grained models are likely to be more complex, containing more correlations than in less drastic coarse-graining. For example, CG descriptions of conjugated polymers require complicated directional interactions. Thus, in highly coarse-grained models, separating the potential of mean force into pairwise site–site interactions will be a worse approximation, and we expect this to be a good area for directly applying GAP-CG.

In this thesis I focused on modelling the *non-bonded* interactions between two and more molecules. However, on a highly coarse-grained scale it turns out that even monomer interactions cannot be separated into the commonly used bond, angle, and dihedral terms. In fact, we have already planned a cooperation with the Voth group in which we will explore the use of GAP for describing the monomer term for individual lipid molecules, combined with their force-matched non-bonded interactions between different molecules.

Adaptive resolution simulations. GAP may prove to be useful in adaptive resolution simulations where part of the system is described at a higher resolution than the rest. An example would be a water-solvated protein with an active binding site and different competing solutes. The binding site and nearby water and solute molecules might require a fully atomistic description, whereas the rest of the protein and solutes that are far away from the binding site do not need such a detailed representation. A significant problem in such simulations is the difference in chemical potential between the CG and the all-atom description of a molecule, due to an incorrect CG description of the thermodynamics. This leads to an artificial drift across the border between CG and all-atom description [142]. So far, simulations need to be corrected manually for this drift, which is a difficult and *ad-hoc* procedure. With the improved representation of both structure and forces in GAP-CG, these manual corrections might no longer be necessary.

Transferability. Another major open question in CG research concerns the *transferability* of CG potentials between different state points. The averaging over microscopic degrees of freedom leads to an entropic component in the CG interactions that inherently depends on the state point, that is, temperature, pressure/density, and so forth. CG potentials that include density or temperature as extended variables have been investigated within the MSCG approach [20–22, 143]. So far, we have not discussed this dependence in detail but, in principle, external parameters such as the temperature could be explicitly added as extended variables to the Gaussian process regression within the GAP approach. This would allow "reuse" of contributions to the potential of mean force that only change slowly with temperature.

A. Appendix

A.1. GAP hyperparameters

hyperparameter	dimer
core potential:	
monomer GAP from grid MF	
and non-bonded short-range core	
covariance cutoff	9 Å
transition width	3 Å
energy scale δ	1 eV
length scale θ	
bonded	3 Å
non-bonded	2 Å
transfer function [Eq. (4.3.2)]	
scaling factor α	4
centre r_0	4 Å
transition width w	2 Å
sparse points n_X	5000
sparse point selection	uniform
noise level (forces) σ_f	0.5 meV/Å

Table A.1: Hyperparameters for the methanol dimer GAP.

hyperparameter	trimer
core potential:	
dimer GAP (Table A.1)	
covariance cutoff	6 Å
transition width	3 Å
energy scale δ	1 eV
length scale θ	
bonded	3 Å
non-bonded	2 Å
sparse points n_X	5000
sparsification method	random
noise level (forces) σ_f	0.5 meV/Å

Table A.2: Hyperparameters for the methanol trimer GAP.

hyperparameter	monomer	dimer
core potential:		
non-bonded short-range core, *cf.* Fig. A.4		
covariance cutoff	—	8 Å
transition width		4.5 Å
energy scale δ	10 eV	1 eV
length scale θ		
bonded	0.3 Å	3 Å
non-bonded	—	1.5 Å
transfer function [Eq. (4.3.2)]		
scaling factor α	—	4
centre r_0	—	4 Å
transition width w	—	2 Å
sparse points n_X	100	2000
sparsification method	uniform	uniform
noise level (forces) σ_f	10 meV/Å	

Table A.3: Hyperparameters for the bulk benzene dimer GAP.

A.2. Pair potentials

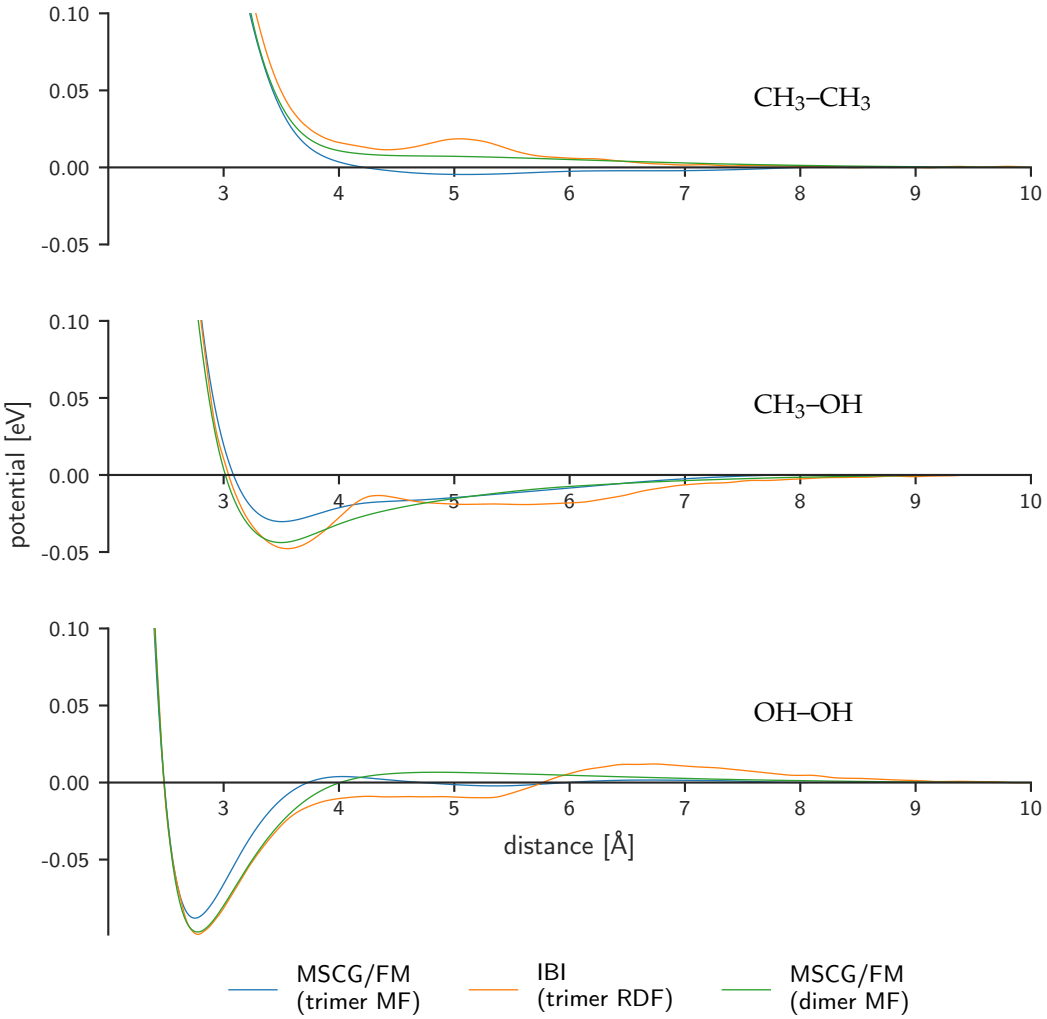

Figure A.1: Methanol trimer pair potentials. From top to bottom: CH_3–CH_3, CH_3–OH and OH–OH interactions. We show MSCG/FM splines obtained from trimer MF (and, for comparison, the corresponding potentials from the dimer system) and the IBI pair potentials that match the trimer site–site distance distributions.

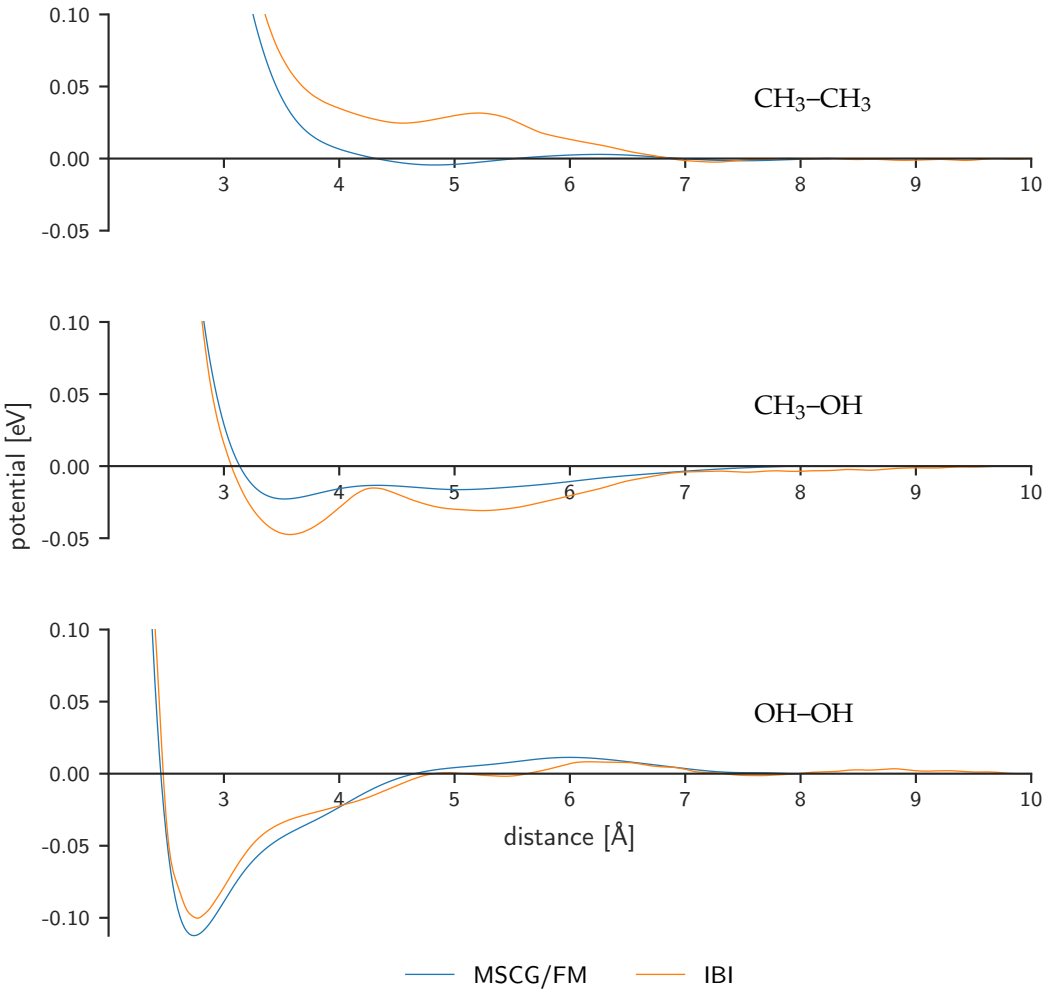

Figure A.2: Methanol tetramer pair potentials. From top to bottom: CH_3–CH_3, CH_3–OH and OH–OH interactions. We show MSCG/FM splines obtained from tetramer MF and the IBI pair potentials that match the tetramer site–site distance distributions.

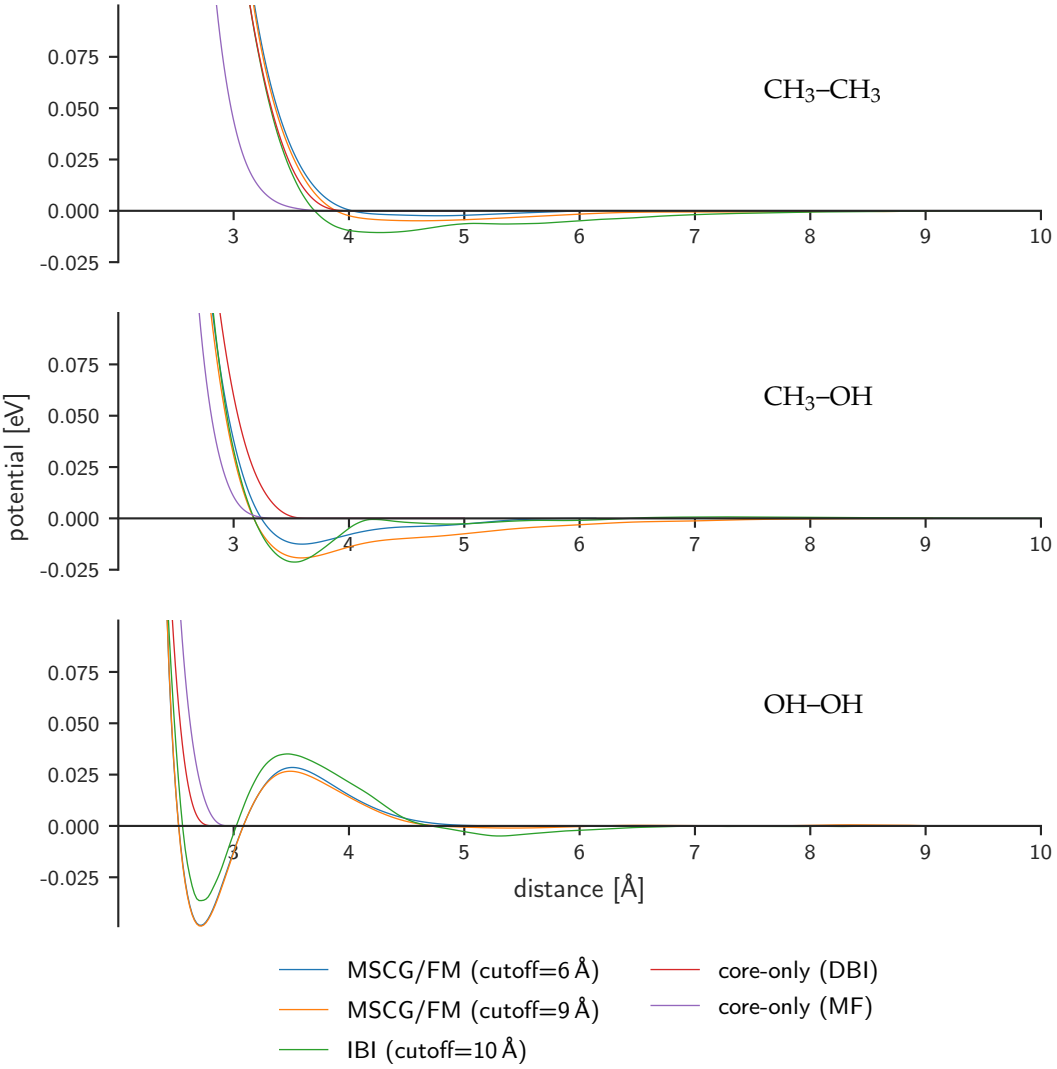

Figure A.3: Bulk methanol pair potentials generated by force-matching (MSCG/FM) and structure-matching (IBI), as well as the two "core-only" short-range potentials.

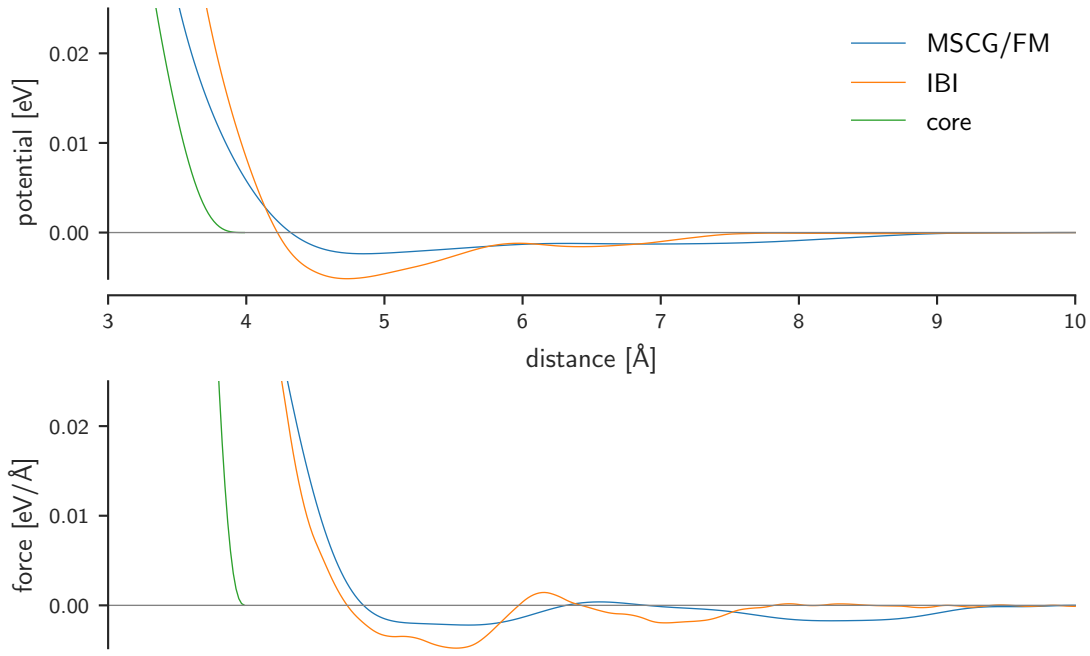

Figure A.4: Bulk benzene pair potentials generated by force-matching (MSCG/FM) and structure-matching (IBI), as well as the short-ranged repulsive core potential.

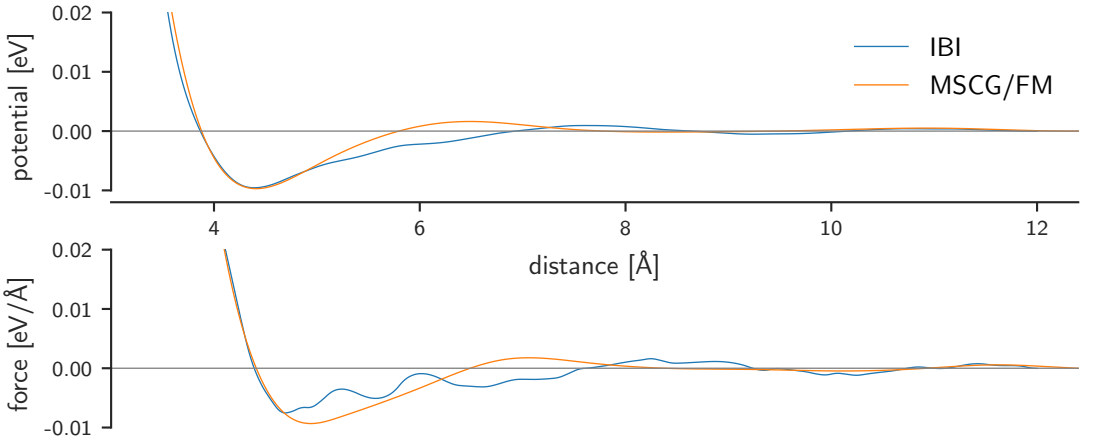

Figure A.5: Pair potentials for water-solvated benzene. The MSCG/FM potential is based on matching the mean forces of the water-solvated dimer and the IBI potential is based on reproducing the site–site RDF of the water-solvated octamer.

A.3. Length scale hyperparameter and transfer function

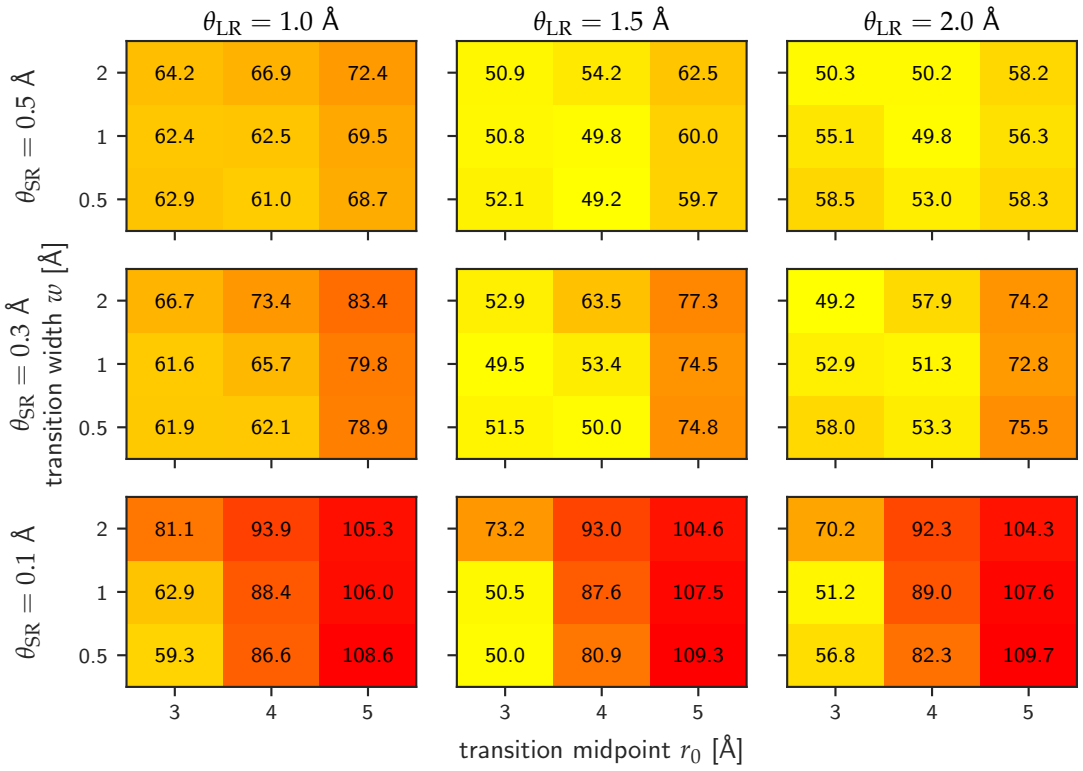

Figure A.6: Influence of length scale hyperparameter and transfer function on GAP force prediction error. Force RMSE in meV/Å for different short- and long-range length scales and for different values of the transition width w and centre r_0 of the transfer function between the two regimes. (All other hyperparameters as given in Table A.3.) If no transfer function is used, there is only a single length scale. The RMSE for different values of this length scale is as follows:

θ/Å	RMSE / meV/Å
0.5	92.4
1.0	64.2
1.5	53.6
2.0	60.0

A.4. Mean forces for benzene dimer in different environments

In Section 6.1.2 we compare the potential of mean force (PMF) for the COM distance between two benzene molecules in different environments: in gas phase, in bulk benzene, and solvated by water. We obtained the mean forces (MF) by Constrained MD using PMFlib. It is possible to use either a single constraint on the distance between the COMs or position constraints on the coordinates of both COMs. The mean force obtained for a distance constraint, $F_{\text{dis}}^{\text{MF}}$, has an additional entropic component due to the rotation of the constraint in space. It is connected to the mean force between the two COMs for position constraints, $F_{\text{pos}}^{\text{MF}}$, by

$$F_{\text{pos}}^{\text{MF}}(r) = F_{\text{dis}}^{\text{MF}}(r) - 2\frac{k_{\text{B}}T}{r},$$

where r is the distance between the COMs. We applied this correction to all mean forces obtained for distance constraints. The resulting mean forces were interpolated by splines which could then be integrated to obtain the PMF. The individual mean force observations, the spline interpolation, and the final PMF are shown in Fig. A.7 on the following page.

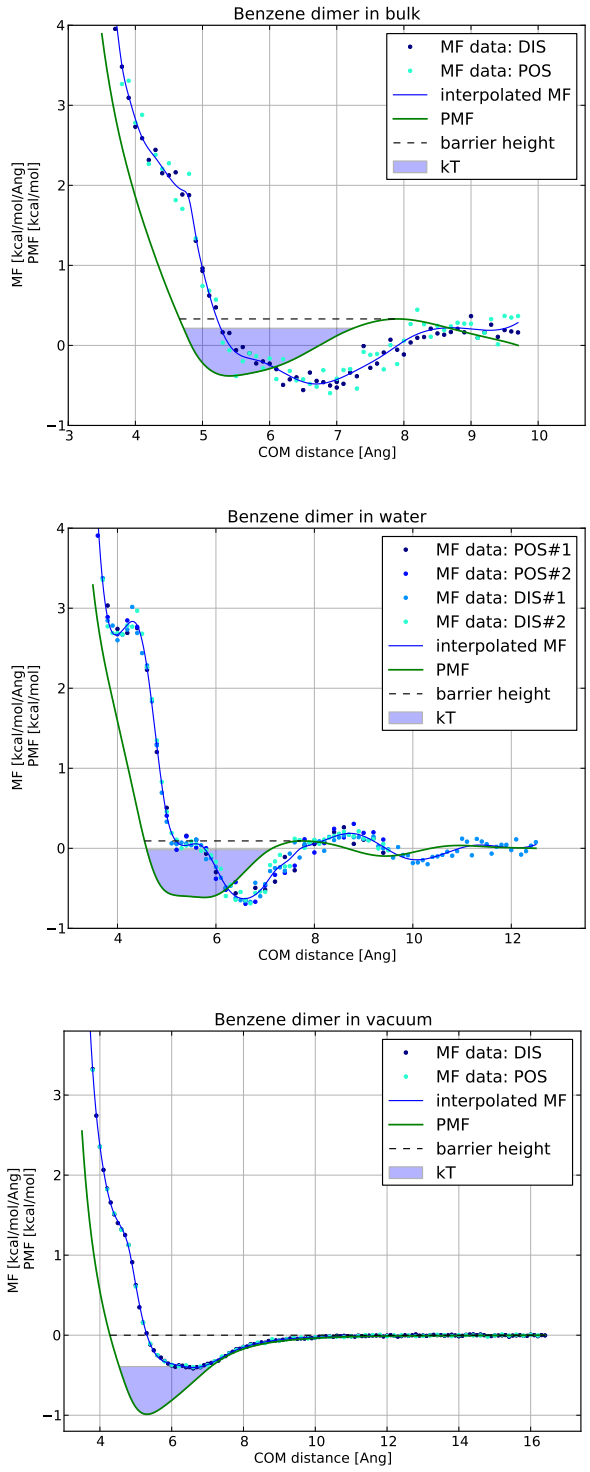

Figure A.7: Benzene in different environments: COM mean forces and PMF.

Bibliography

[1] W. F. van Gunsteren and H. J. C. Berendsen. "Computer simulation of molecular dynamics: Methodology, applications, and perspectives in chemistry". *Angewandte Chemie International Edition in English* **29**, 992 (1990).

[2] J. R. Kermode, T. Albaret, D. Sherman, N. Bernstein, P. Gumbsch, M. C. Payne, G. Csányi, and A. De Vita. "Low-speed fracture instabilities in a brittle crystal". *Nature* **455**, 1224 (2008).

[3] M. Born and R. Oppenheimer. "Zur Quantentheorie der Molekeln". *Annalen der Physik* **389**, 457 (1927).

[4] D. Frenkel and B. Smit. *Understanding Molecular Simulation*. Academic Press, 2001.

[5] R. P. Feynman. "Forces in molecules". *Phys. Rev.* **56**, 340 (1939).

[6] D. Ceperley and B. Alder. "Quantum Monte Carlo". *Science* **231**, 555 (1986).

[7] R. G. Parr and W. Yang. *Density–Functional Theory of Atoms and Molecules, International Series of Monographs on Chemistry, vol. 16*. Oxford University Press, 1989.

[8] N. D. M. Hine, M. Robinson, P. D. Haynes, C.-K. Skylaris, M. C. Payne, and A. A. Mostofi. "Accurate ionic forces and geometry optimization in linear-scaling density-functional theory with local orbitals". *Phys. Rev. B* **83**, 195102 (2011).

[9] D. E. Shaw, M. M. Deneroff, R. O. Dror, J. S. Kuskin, R. H. Larson, J. K. Salmon, C. Young, B. Batson, K. J. Bowers, J. C. Chao, M. P. Eastwood, J. Gagliardo, J. P. Grossman, C. R. Ho, D. J. Ierardi, I. Kolossváry, J. L. Klepeis, T. Layman, C. McLeavey, M. A. Moraes, R. Mueller, E. C. Priest, Y. Shan, J. Spengler, M. Theobald, B. Towles, and S. C. Wang. "Anton, a special-purpose machine for molecular dynamics simulation". *Communications of the ACM* **51**, 91 (2008).

[10] K. Lindorff-Larsen, P. Maragakis, S. Piana, and D. E. Shaw. "Picosecond to millisecond structural dynamics in human ubiquitin". *The Journal of Physical Chemistry B* (2016).

[11] G. Zhao, J. R. Perilla, E. L. Yufenyuy, X. Meng, B. Chen, J. Ning, J. Ahn, A. M. Gronenborn, K. Schulten, C. Aiken, and P. Zhang. "Mature HIV-1 capsid structure by cryo-electron microscopy and all-atom molecular dynamics". *Nature* **497**, 643 (2013).

[12] T. Reddy, D. Shorthouse, D. Parton, E. Jefferys, P. Fowler, M. Chavent, M. Baaden, and M. Sansom. "Nothing to sneeze at: A dynamic and integrative computational model of an influenza a virion". *Structure* **23**, 584 (2015).

[13] M. L. Klein and W. Shinoda. "Large-scale molecular dynamics simulations of self-assembling systems". *Science* **321**, 798 (2008).

[14] G. A. Voth (Ed.). *Coarse-Graining of Condensed Phase and Biomolecular Systems*. CRC Press, 2008.

[15] R. M. Jendrejack, E. T. Dimalanta, D. C. Schwartz, M. D. Graham, and J. J. de Pablo. "DNA dynamics in a microchannel". *Phys. Rev. Lett.* **91**, 038102 (2003).

[16] S. J. Marrink, H. J. Risselada, S. Yefimov, D. P. Tieleman, and A. H. de Vries. "The MARTINI force field: Coarse grained model for biomolecular simulations". *The Journal of Physical Chemistry B* **111**, 7812 (2007).

[17] H. I. Ingólfsson, C. A. Lopez, J. J. Uusitalo, D. H. de Jong, S. M. Gopal, X. Periole, and S. J. Marrink. "The power of coarse graining in biomolecular simulations". *Wiley Interdisciplinary Reviews: Computational Molecular Science* **4**, 225 (2014).

[18] M. Praprotnik, L. D. Site, and K. Kremer. "Multiscale simulation of soft matter: From scale bridging to adaptive resolution". *Annual Review of Physical Chemistry* **59**, 545 (2008).

[19] P. Carbone, H. A. K. Varzaneh, X. Chen, and F. Müller-Plathe. "Transferability of coarse-grained force fields: The polymer case". *The Journal of Chemical Physics* **128**, 064904 (2008).

[20] V. Krishna, W. G. Noid, and G. A. Voth. "The multiscale coarse-graining method. IV. transferring coarse-grained potentials between temperatures". *The Journal of Chemical Physics* **131**, 024103 (2009).

[21] S. Izvekov, P. W. Chung, and B. M. Rice. "The multiscale coarse-graining method: Assessing its accuracy and introducing density dependent coarse-grain potentials". *The Journal of Chemical Physics* **133**, 064109 (2010).

[22] J. W. Mullinax and W. G. Noid. "Extended ensemble approach for deriving transferable coarse-grained potentials". *The Journal of Chemical Physics* **131**, 104110 (2009).

[23] S. Izvekov and G. A. Voth. "Modeling real dynamics in the coarse-grained representation of condensed phase systems". *The Journal of Chemical Physics* **125**, 151101 (2006).

[24] A. A. Louis. "Beware of density dependent pair potentials". *Journal of Physics: Condensed Matter* **14**, 9187 (2002).

[25] M. E. Johnson, T. Head-Gordon, and A. A. Louis. "Representability problems for coarse-grained water potentials". *The Journal of Chemical Physics* **126**, 144509 (2007).

[26] L. Larini, L. Lu, and G. A. Voth. "The multiscale coarse-graining method. VI. implementation of three-body coarse-grained potentials". *The Journal of Chemical Physics* **132**, 164107 (2010).

[27] A. Das and H. C. Andersen. "The multiscale coarse-graining method. IX. a general method for construction of three body coarse-grained force fields" .

[28] A. Bartók-Pártay. *Gaussian Approximation Potential: an interatomic potential derived from first principles Quantum Mechanics*. Ph.D. thesis, University of Cambridge (2010).

[29] A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi. "Gaussian Approximation Potentials: The accuracy of Quantum Mechanics, without the electrons". *Phys. Rev. Lett.* **104**, 136403 (2010).

[30] A. P. Bartók and G. Csányi. "Gaussian approximation potentials: A brief tutorial introduction". *International Journal of Quantum Chemistry* **115**, 1051 (2015).

[31] D. J. C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.

[32] C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.

[33] W. C. Swope, H. C. Andersen, P. H. Berens, and K. R. Wilson. "A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: Application to small water clusters". *The Journal of Chemical Physics* **76**, 637 (1982).

[34] M. P. Allen and D. J. Tildesley. *Computer simulation of liquids*. Oxford University Press, 1989.

[35] W. D. Cornell, P. Cieplak, C. I. Bayly, I. R. Gould, K. M. Merz, D. M. Ferguson, D. C. Spellmeyer, T. Fox, J. W. Caldwell, and P. A. Kollman. "A second generation force field for the simulation of proteins, nucleic acids, and organic molecules". *Journal of the American Chemical Society* **117**, 5179 (1995).

[36] V. Hornak, R. Abel, A. Okur, B. Strockbine, A. Roitberg, and C. Simmerling. "Comparison of multiple Amber force fields and development of improved protein backbone parameters". *Proteins: Structure, Function, and Bioinformatics* **65**, 712 (2006).

[37] W. L. Jorgensen. *OPLS Force Fields*. John Wiley & Sons, Ltd, 2002.

[38] A. D. MacKerell, N. Banavali, and N. Foloppe. "Development and current status of the CHARMM force field for nucleic acids". *Biopolymers* **56**, 257 (2000).

[39] W. F. van Gunsteren, X. Daura, and A. E. Mark. *GROMOS Force Field*. John Wiley & Sons, Ltd, 2002.

[40] N. Schmid, A. P. Eichenberger, A. Choutko, S. Riniker, M. Winger, A. E. Mark, and W. F. van Gunsteren. "Definition and testing of the GROMOS force-field versions 54A7 and 54B7". *European Biophysics Journal* **40**, 843 (2011).

[41] T. A. Halgren and W. Damm. "Polarizable force fields". *Current Opinion in Structural Biology* **11**, 236 (2001).

[42] J. W. Ponder, C. Wu, P. Ren, V. S. Pande, J. D. Chodera, M. J. Schnieders, I. Haque, D. L. Mobley, D. S. Lambrecht, J. Robert A. DiStasio, M. Head-Gordon, G. N. I. Clark, M. E. Johnson, and T. Head-Gordon. "Current status of the AMOEBA polarizable force field". *The Journal of Physical Chemistry B* **114**, 2549 (2010).

[43] G. A. Cisneros, K. T. Wikfeldt, L. Ojamäe, J. Lu, Y. Xu, H. Torabifard, A. P. Bartók, G. Csányi, V. Molinero, and F. Paesani. "Modeling molecular interactions in water: From pairwise to many-body potential energy functions". *Chemical Reviews* **116**, 7501 (2016).

[44] A. Warshel and R. M. Weiss. "An empirical valence bond approach for comparing reactions in solutions and in enzymes". *Journal of the American Chemical Society* **102**, 6218 (1980).

[45] A. C. T. van Duin, S. Dasgupta, F. Lorant, and W. A. Goddard III. "ReaxFF: A reactive force field for hydrocarbons". *The Journal of Physical Chemistry A* **105**, 9396 (2001).

[46] D. W. Hofmann, L. Kuleshova, and B. D'Aguanno. "A new reactive potential for the molecular dynamics simulation of liquid water". *Chemical Physics Letters* **448**, 138 (2007).

[47] J. L. Klepeis, K. Lindorff-Larsen, R. O. Dror, and D. E. Shaw. "Long-timescale molecular dynamics simulations of protein structure and function". *Current Opinion in Structural Biology* **19**, 120 (2009).

[48] H. D. Herce and A. E. Garcia. "Molecular dynamics simulations suggest a mechanism for translocation of the HIV-1 TAT peptide across lipid membranes". *Proceedings of the National Academy of Sciences* **104**, 20805 (2007).

[49] D. Tieleman, S. Marrink, and H. Berendsen. "A computer perspective of membranes: molecular dynamics studies of lipid bilayer systems". *Biochimica et Biophysica Acta (BBA) - Reviews on Biomembranes* **1331**, 235 (1997).

[50] T. Darden, D. York, and L. Pedersen. "Particle mesh Ewald: An $N \cdot N \log(N)$ method for Ewald sums in large systems". *The Journal of Chemical Physics* **98**, 10089 (1993).

[51] J. N. Israelachvili. *Intermolecular and surface forces*. Academic Press, 2011.

[52] A. Onufriev. "Implicit solvent models in molecular dynamics simulations: A brief overview". *Annual Reports in Computational Chemistry* **4**, 125 (2008).

[53] D. Chandler. "Interfaces and the driving force of hydrophobic assembly". *Nature* **437**, 640 (2005).

[54] I. Daidone, M. B. Ulmschneider, A. Di Nola, A. Amadei, and J. C. Smith. "Dehydration-driven solvent exposure of hydrophobic surfaces as a driving force in peptide folding". *Proceedings of the National Academy of Sciences* **104**, 15230 (2007).

[55] P. H. Hünenberger. "Thermostat algorithms for molecular dynamics simulations". C. Holm and K. Kremer (Eds.), "Advanced Computer Simulation", *Advances in Polymer Science*, vol. 173, 130–130. Springer, 2005.

[56] F. Legoll, M. Luskin, and R. Moeckel. "Non-ergodicity of the Nosé–Hoover thermostatted harmonic oscillator". *Archive for Rational Mechanics and Analysis* **184**, 449 (2007).

[57] B. Leimkuhler, E. Noorizadeh, and F. Theil. "A gentle stochastic thermostat for molecular dynamics". *Journal of Statistical Physics* **135**, 261 (2009).

[58] R. Kubo. "The fluctuation-dissipation theorem". *Reports on Progress in Physics* **29**, 255 (1966).

[59] J. A. Izaguirre, D. P. Catarello, J. M. Wozniak, and R. D. Skeel. "Langevin stabilization of molecular dynamics". *The Journal of Chemical Physics* **114**, 2090 (2001).

[60] A. Szarecka, R. P. White, and H. Meirovitch. "Absolute entropy and free energy of fluids using the hypothetical scanning method. I. calculation of transition probabilities from local grand canonical partition functions". *The Journal of Chemical Physics* **119**, 12084 (2003).

[61] J. D. Chodera, D. L. Mobley, M. R. Shirts, R. W. Dixon, K. Branson, and V. S. Pande. "Alchemical free energy methods for drug discovery: progress and challenges". *Current Opinion in Structural Biology* **21**, 150 (2011).

[62] C. Chipot and A. Pohorille. *Free Energy Calculations*. Springer, 2007.

[63] P. Kulhánek, L. Mones, Z. Střelcová, I. Simon, M. Fuxreiter, and J. Koča. "PMFLib — a toolkit for free energy calculations". "8th Discussions in Structural Molecular Biology, Nové Hrady, Czech Republic", 2011.

[64] G. Ramachandran, C. Ramakrishnan, and V. Sasisekharan. "Stereochemistry of polypeptide chain configurations". *Journal of Molecular Biology* **7**, 95 (1963).

[65] T. Stecher, N. Bernstein, and G. Csányi. "Free energy surface reconstruction from Umbrella samples using Gaussian Process Regression". *Journal of Chemical Theory and Computation* **10**, 4079 (2014).

[66] L. Mones, N. Bernstein, and G. Csányi. "Exploration, sampling and reconstruction of free energy surfaces with Gaussian process regression". *Journal of Chemical Theory and Computation* (2016).

[67] J. Kästner. "Umbrella sampling". *Wiley Interdisciplinary Reviews: Computational Molecular Science* **1**, 932 (2011).

[68] J. Kästner and W. Thiel. "Bridging the gap between thermodynamic integration and umbrella sampling provides a novel analysis method: "Umbrella integration"". *The Journal of Chemical Physics* **123**, 144104 (2005).

[69] E. Carter, G. Ciccotti, J. T. Hynes, and R. Kapral. "Constrained reaction coordinate dynamics for the simulation of rare events". *Chemical Physics Letters* **156**, 472 (1989).

[70] M. Sprik and G. Ciccotti. "Free energy from constrained molecular dynamics". *The Journal of Chemical Physics* **109**, 7737 (1998).

[71] E. Darve and A. Pohorille. "Calculating free energies using average force". *The Journal of Chemical Physics* **115**, 9169 (2001).

[72] J. Hénin and C. Chipot. "Overcoming free energy barriers using unconstrained molecular dynamics simulations". *The Journal of Chemical Physics* **121**, 2904 (2004).

[73] A. Laio and M. Parrinello. "Escaping free-energy minima". *Proceedings of the National Academy of Sciences* **99**, 12562 (2002).

[74] A. Barducci, G. Bussi, and M. Parrinello. "Well-tempered metadynamics: A smoothly converging and tunable free-energy method". *Phys. Rev. Lett.* **100**, 020603 (2008).

[75] A. Laio and F. L. Gervasio. "Metadynamics: a method to simulate rare events and reconstruct the free energy in biophysics, chemistry and material science". *Reports on Progress in Physics* **71**, 126601 (2008).

[76] J. Kästner and W. Thiel. "Analysis of the statistical error in umbrella sampling simulations by umbrella integration". *The Journal of Chemical Physics* **124**, 234106 (2006).

[77] A. V. Sinitskiy, M. G. Saunders, and G. A. Voth. "Optimal number of coarse-grained sites in different components of large biomolecular complexes". *The Journal of Physical Chemistry B* **116**, 8363 (2012).

[78] N. Guttenberg, J. F. Dama, M. G. Saunders, G. A. Voth, J. Weare, and A. R. Dinner. "Minimizing memory as an objective for coarse-graining". *The Journal of Chemical Physics* **138**, 094111 (2013).

[79] T. T. Foley, M. S. Shell, and W. G. Noid. "The impact of resolution upon entropy and information in coarse-grained models". *The Journal of Chemical Physics* **143**, 243104 (2015).

[80] W. G. Noid. "Perspective: Coarse-grained models for biomolecular systems". *The Journal of Chemical Physics* **139**, 090901 (2013).

[81] S. Riniker, J. R. Allison, and W. F. van Gunsteren. "On developing coarse-grained models for biomolecular simulation: a review". *Phys. Chem. Chem. Phys.* **14**, 12423 (2012).

[82] J. F. Rudzinski and W. G. Noid. "Investigation of coarse-grained mappings via an iterative generalized Yvon–Born–Green method". *The Journal of Physical Chemistry B* **118**, 8295 (2014).

[83] A. Das, L. Lu, H. C. Andersen, and G. A. Voth. "The multiscale coarse-graining method. X. improved algorithms for constructing coarse-grained potentials for molecular systems". *The Journal of Chemical Physics* **136**, 194115 (2012).

[84] J. F. Rudzinski and W. G. Noid. "Coarse-graining entropy, forces, and structures". *The Journal of Chemical Physics* **135**, 214101 (2011).

[85] T. E. Ouldridge. *Coarse-grained modelling of DNA and DNA self-assembly*. Ph.D. thesis, University of Oxford (2011).

[86] T. E. Ouldridge, A. A. Louis, and J. P. K. Doye. "Structural, mechanical, and thermodynamic properties of a coarse-grained DNA model". *The Journal of Chemical Physics* **134**, 085101 (2011).

[87] J. Comer, J. C. Gumbart, J. Hénin, T. Lelièvre, A. Pohorille, and C. Chipot. "The adaptive biasing force method: Everything you always wanted to know but were afraid to ask". *The Journal of Physical Chemistry B* **119**, 1129 (2015).

[88] A. Srivastava and G. A. Voth. "Hybrid approach for highly coarse-grained lipid bilayer models". *Journal of Chemical Theory and Computation* **9**, 750 (2013).

[89] J. J. de Pablo, B. Jones, C. L. Kovacs, V. Ozolins, and A. P. Ramirez. "The materials genome initiative, the interplay of experiment, theory and computation". *Current Opinion in Solid State and Materials Science* **18**, 99 (2014).

[90] W. G. Noid, J.-W. Chu, G. S. Ayton, V. Krishna, S. Izvekov, G. A. Voth, A. Das, and H. C. Andersen. "The multiscale coarse-graining method. I. a rigorous bridge between atomistic and coarse-grained models". *The Journal of Chemical Physics* **128**, 244114 (2008).

[91] W. G. Noid. *Systematic Methods for Structurally Consistent Coarse-Grained Models*, 487–531. Humana Press, Totowa, NJ, 2013.

[92] R. Henderson. "A uniqueness theorem for fluid pair correlation functions". *Physics Letters A* **49**, 197 (1974).

[93] E. Brini, E. A. Algaer, P. Ganguly, C. Li, F. Rodriguez-Ropero, and N. F. A. van der Vegt. "Systematic coarse-graining methods for soft matter simulations — a review". *Soft Matter* **9**, 2108 (2013).

[94] V. Rühle, C. Junghans, A. Lukyanov, K. Kremer, and D. Andrienko. "Versatile object-oriented toolkit for coarse-graining applications". *Journal of Chemical Theory and Computation* **5**, 3211 (2009).

[95] W. Tschöp, K. Kremer, J. Batoulis, T. Bürger, and O. Hahn. "Simulation of polymer melts. I. coarse-graining procedure for polycarbonates". *Acta Polymerica* **49**, 61 (1998).

[96] D. Chandler. *Introduction to modern statistical mechanics*. Oxford University Press, 1987.

[97] D. Reith, M. Pütz, and F. Müller-Plathe. "Deriving effective mesoscale potentials from atomistic simulations". *Journal of Computational Chemistry* **24**, 1624 (2003).

[98] C. Peter. "Coarse grained models for multiscale simulations of biomolecular systems". J. Grotendorst (Ed.), "Hierarchical Methods for Dynamics in Complex Molecular Systems", vol. 10, 195. Forschungszentrum Jülich, 2012.

[99] A. P. Lyubartsev and A. Laaksonen. "Calculation of effective interaction potentials from radial distribution functions: A reverse Monte Carlo approach". *Phys. Rev. E* **52**, 3730 (1995).

[100] A. Soper. "Empirical potential Monte Carlo simulation of fluid structure". *Chemical Physics* **202**, 295 (1996).

[101] A. P. Lyubartsev, M. Karttunen, I. Vattulainen, and A. Laaksonen. "On coarse-graining by the Inverse Monte Carlo method: Dissipative Particle Dynamics simulations made to a precise tool in soft matter modeling". *Soft Materials* **1**, 121 (2002).

[102] M. S. Shell. "The relative entropy is fundamental to multiscale and inverse thermodynamic problems". *The Journal of Chemical Physics* **129**, 144108 (2008).

[103] S. Y. Mashayak, M. N. Jochum, K. Koschke, N. R. Aluru, V. Rühle, and C. Junghans. "Relative entropy and optimization-driven coarse-graining methods in VOTCA". *PLoS ONE* **10**, 1 (2015).

[104] A. Chaimovich and M. S. Shell. "Coarse-graining errors and numerical optimization using a relative entropy framework". *The Journal of Chemical Physics* **134**, 094112 (2011).

[105] F. Ercolessi and J. B. Adams. "Interatomic potentials from first-principles calculations: The force-matching method". *EPL (Europhysics Letters)* **26**, 583 (1994).

[106] S. Izvekov and G. A. Voth. "A multiscale coarse-graining method for biomolecular systems". *The Journal of Physical Chemistry B* **109**, 2469 (2005).

[107] S. Izvekov and G. Voth. "Multiscale coarse graining of liquid-state systems". *The Journal of Chemical Physics* **123**, 134105 (2005).

[108] W. G. Noid, P. Liu, Y. Wang, J.-W. Chu, G. S. Ayton, S. Izvekov, H. C. Andersen, and G. A. Voth. "The multiscale coarse-graining method. II. numerical implementation for coarse-grained molecular models". *The Journal of Chemical Physics* **128**, 244115 (2008).

[109] T. Lelièvre, G. Stoltz, and M. Rousset. *Free energy computations: A mathematical perspective*. World Scientific, 2010.

[110] M. J. Mills and P. L. Popelier. "Intramolecular polarisable multipolar electrostatics from the machine learning method Kriging". *Computational and Theoretical Chemistry* **975**, 42 (2011).

[111] M. Mills and P. Popelier. "Polarisable multipolar electrostatics from the machine learning method Kriging: an application to alanine". *Theoretical Chemistry Accounts: Theory, Computation, and Modeling (Theoretica Chimica Acta)* **131**, 1 (2012).

[112] J. Behler. "Representing potential energy surfaces by high-dimensional neural network potentials". *Journal of Physics: Condensed Matter* **26**, 183001 (2014).

[113] F. H. Stillinger and T. A. Weber. "Computer simulation of local order in condensed phases of silicon". *Phys. Rev. B* **31**, 5262 (1985).

[114] M. S. Daw, S. M. Foiles, and M. I. Baskes. "The embedded-atom method: a review of theory and applications". *Materials Science Reports* **9**, 251 (1993).

[115] M. W. Finnis and J. E. Sinclair. "A simple empirical N-body potential for transition metals". *Philosophical Magazine A* **50**, 45 (1984).

[116] L. Paramonov, M. Burke, and S. Yaliraki. "Coarse-grained dynamics of anisotropic systems". G. Voth (Ed.), "Coarse-Graining of Condensed Phase and Biomolecular Systems", 59–68. CRC Press, Boca Raton, FL, 2009.

[117] J. M. Phillips and S. Venkatasubramanian. "A Gentle Introduction to the Kernel Distance". *ArXiv e-prints* (2011).

[118] A. P. Bartók, R. Kondor, and G. Csányi. "On representing chemical environments". *Phys. Rev. B* **87**, 184115 (2013).

[119] S. Barnett and S. Barnett. *Matrix methods for engineers and scientists*. McGraw-Hill, 1979.

[120] P. Liu, Q. Shi, H. Daumé, and G. A. Voth. "A bayesian statistics approach to multiscale coarse graining". *The Journal of Chemical Physics* **129**, 214114 (2008).

[121] E. Solak, R. Murray-Smith, W. Leithead, D. Leith, and C. Rasmussen. "Derivative observations in Gaussian Process models of dynamic systems" (2003).

[122] C. Rasmussen. "Gaussian Processes to speed up Hybrid Monte Carlo for expensive Bayesian integrals". J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith, and M. West (Eds.), "Bayesian Statistics 7", 651–659. 2003.

[123] J. Quiñonero-Candela and C. E. Rasmussen. "A unifying view of sparse approximate Gaussian process regression". *Journal of Machine Learning Research* **6**, 1939 (2005).

[124] E. Snelson and Z. Ghahramani. "Sparse Gaussian processes using pseudo-inputs". "Advances in neural information processing systems", 1257–1264. 2005.

[125] Z. Cao, J. F. Dama, L. Lu, and G. A. Voth. "Solvent free ionic solution models from multiscale coarse-graining". *Journal of Chemical Theory and Computation* **9**, 172 (2013).

[126] M. W. Mahoney and P. Drineas. "CUR matrix decompositions for improved data analysis". *Proceedings of the National Academy of Sciences* **106**, 697 (2009).

[127] J. Schreiter, D. Nguyen-Tuong, and M. Toussaint. "Efficient sparsification for Gaussian process regression". *Neurocomputing* **192**, 29 (2016).

[128] C. Rasmussen and Z. Ghahramani. "Occam's razor". *Advances in Neural Information Processing Systems* 294 (2001).

[129] M. Gibbs. *Bayesian Gaussian processes for classification and regression*. Ph.D. thesis, University of Cambridge (1997).

[130] Z. Cao and G. A. Voth. "The multiscale coarse-graining method. XI. accurate interactions based on the centers of charge of coarse-grained sites". *The Journal of Chemical Physics* **143**, 243116 (2015).

[131] D. Case *et al.* "Amber 12" (2012). University of California, San Francisco.

[132] M. J. Frisch *et al.* "Gaussian 09 Revision E.01" (2009). Gaussian Inc. Wallingford CT.

[133] G. Csányi *et al.* "QUIP" (2006–2016). University of Cambridge.

[134] S. Plimpton. "Fast parallel algorithms for short-range molecular dynamics". *Journal of Computational Physics* **117**, 1 (1995).

[135] J. F. Rudzinski and W. G. Noid. "The role of many-body correlations in determining potentials for coarse-grained models of equilibrium structure". *The Journal of Physical Chemistry B* **116**, 8621 (2012).

[136] W. L. Jorgensen, J. Chandrasekhar, J. D. Madura, R. W. Impey, and M. L. Klein. "Comparison of simple potential functions for simulating liquid water". *The Journal of Chemical Physics* **79**, 926 (1983).

[137] W. C. Still, A. Tempczyk, R. C. Hawley, and T. Hendrickson. "Semianalytical treatment of solvation for molecular mechanics and dynamics". *Journal of the American Chemical Society* **112**, 6127 (1990).

[138] A. Villa, C. Peter, and N. F. A. van der Vegt. "Transferability of nonbonded interaction potentials for coarse-grained simulations: Benzene in water". *Journal of Chemical Theory and Computation* **6**, 2434 (2010).

[139] G. D. Hawkins, C. J. Cramer, and D. G. Truhlar. "Pairwise solute descreening of solute charges from a dielectric medium". *Chemical Physics Letters* **246**, 122 (1995).

[140] G. D. Hawkins, C. J. Cramer, and D. G. Truhlar. "Parametrized models of aqueous free energies of solvation based on pairwise descreening of solute atomic charges from a dielectric medium". *The Journal of Physical Chemistry* **100**, 19824 (1996).

[141] J. M. A. Grime and G. A. Voth. "Highly scalable and memory efficient ultra-coarse-grained molecular dynamics simulations". *Journal of Chemical Theory and Computation* **10**, 423 (2014).

[142] S. Fritsch, S. Poblete, C. Junghans, G. Ciccotti, L. Delle Site, and K. Kremer. "Adaptive resolution molecular dynamics simulation through coupling to an internal particle reservoir". *Phys. Rev. Lett.* **108**, 170602 (2012).

[143] S. Izvekov, P. W. Chung, and B. M. Rice. "Particle-based multiscale coarse graining with density-dependent potentials: Application to molecular crystals (hexahydro-1,3,5-trinitro-s-triazine)". *The Journal of Chemical Physics* **135**, 044112 (2011).